

How can one test if a binary sequence is exchangeable? Fork-convex hulls, supermartingales, and e-processes

Aaditya Ramdas¹, Johannes Ruf², Martin Larsson³, Wouter M. Koolen⁴

¹ Departments of Statistics and ML, Carnegie Mellon University

² Department of Mathematics, London School of Economics

³ Department of Mathematics, Carnegie Mellon University

⁴ Machine Learning Group, CWI Amsterdam

{aramdas,martinl}@andrew.cmu.edu
wmkoolen@cw.nl, j.ruf@lse.ac.uk

July 1, 2021

Abstract

Suppose we observe an infinite series of coin flips X_1, X_2, \dots , and wish to sequentially test the null that these binary random variables are exchangeable. Nonnegative supermartingales (NSMs) are a workhorse of sequential inference, but we prove that they are powerless for this problem. First, utilizing a geometric concept called fork-convexity (a sequential analog of convexity), we show that any process that is an NSM under two distributions, is also necessarily an NSM under their “fork-convex hull”. Second, we demonstrate that the fork-convex hull of the exchangeable null consists of all possible laws over binary sequences; this implies that any NSM under exchangeability is necessarily nonincreasing, hence always yields a powerless test for any alternative. Since testing arbitrary deviations from exchangeability is information theoretically impossible, we focus on Markovian alternatives. We combine ideas from universal inference and the method of mixtures to derive a “safe e-process”, which is a nonnegative process with expectation at most one under the null at any stopping time, and is upper bounded by a martingale, but is not itself an NSM. This in turn yields a level α sequential test that is consistent; regret bounds from universal coding also demonstrate rate-optimal power. We present ways to extend these results to any finite alphabet and to Markovian alternatives of any order using a “double mixture” approach. We provide a wide array of simulations, and give general approaches based on betting for unstructured or ill-specified alternatives. Finally, inspired by Shafer, Vovk, and Ville, we provide game-theoretic interpretations of our e-processes and pathwise results.

Keywords: Anytime-valid sequential inference; betting; calibrator; composite Snell envelope; de Finetti mixing; fork-convexity; game-theoretic probability; Jeffreys’ prior; method of mixtures; nonnegative supermartingale; optional stopping; regret bound; safe e-process; universal coding; Ville’s theorem.

Contents

1	Testing exchangeability	3
1.1	Safe e-processes and anytime-valid p-processes	3
1.2	Convex hulls and de Finetti's theorem	4
1.3	Time-homogeneous Markovian alternatives	5
1.4	Wald's sequential likelihood (or probability) ratio test	6
1.5	Nonnegative supermartingales (NSMs) and an impossibility result	6
1.6	A powerful \mathcal{Q} -safe e-process that is not a \mathcal{Q} -NSM	7
1.7	A summary of this paper's contributions	8
2	Jeffreys' mixture meets maximum likelihood	8
2.1	Dealing with the composite null via maximum likelihood	8
2.2	Dealing with the composite alternative using a Jeffreys' mixture	9
2.3	Combining Jeffreys' mixture with maximum likelihood	10
2.4	A pathwise theorem on the power of (R_t)	10
2.5	Extension to higher-order Markovian alternatives and context trees	12
2.6	Double-mixtures for a countable sequence of alternatives	13
2.7	Handling generic alternatives via non-anticipating likelihoods and betting	13
3	Fork-convexity and \mathcal{Q}-Snell envelopes	14
3.1	A sequential analog of convexity	14
3.2	No power against fork-convex hulls	15
3.3	Composite Snell envelopes	16
3.4	The inadequacy of NSMs for testing exchangeability	17
4	Numerical simulations and some extensions	18
4.1	Sanity check 1: no power against an i.i.d. Bernoulli sequence	19
4.2	Sanity check 2: power against a first-order Markov alternative	20
4.3	Power against a time-varying Markov alternative	20
4.4	Power against a change point alternative	21
4.5	Power on a real-world data set	23
4.6	An equivalent viewpoint based on confidence sequences	23
4.7	Calibrated p-processes and adjusted e-processes for not losing capital	24
5	Summary and discussion	25
5.1	From convex hulls to fork-convex hulls	26
5.2	The powerlessness of test supermartingales, and a powerful e-process	26
5.3	Vovk's approach based on conformal prediction	26
5.4	A game-theoretic protocol and reinterpretation of Theorem 7	27
6	References	28
A	Additional technical concepts and definitions	31
A.1	Reference measures and local absolute continuity	31
A.2	Essential supremum	32
B	Omitted proofs	32

1 Testing exchangeability

Suppose we observe a sequence of binary coin flips X_1, X_2, \dots one at a time. Consider the fundamental problem of testing if our data $(X_t)_{t \geq 1}$ is an exchangeable sequence:

$$H_0 : X_1, X_2, \dots \text{ are exchangeable.}$$

For the unfamiliar reader, $(X_t)_{t \geq 1}$ is exchangeable if and only if for every t and every permutation σ of the first t indices, (X_1, \dots, X_t) has the same distribution as $(X_{\sigma(1)}, \dots, X_{\sigma(t)})$. If we find enough evidence against the null, we would like to stop collecting data and reject the null as soon as possible. Let $\overline{\mathcal{Q}}$ represent the set of all exchangeable distributions over infinite binary sequences. Then, we can rephrase the null as H_0 : the data are generated from some $\mathbf{Q} \in \overline{\mathcal{Q}}$.

Let $(\mathcal{F}_t)_{t \geq 0}$ represent the canonical filtration of the data, where \mathcal{F}_0 is the trivial sigma algebra and $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$. All martingale statements in this paper will implicitly refer to this canonical filtration. A level α sequential test for $\overline{\mathcal{Q}}$ (that is, for H_0) is any stopping time τ_α such that

$$\sup_{\mathbf{Q} \in \overline{\mathcal{Q}}} \mathbf{Q}(\tau_\alpha < \infty) \leq \alpha, \tag{1}$$

meaning that with probability $1 - \alpha$, we never stop under the null. In this paper, we will design powerful sequential tests for exchangeability, and prove their consistency, and in some cases rate-optimality, against Markovian alternatives.

Despite nonnegative supermartingales (NSMs) being a central object in sequential testing [12, 13], the nontriviality of our contribution stems from the following: Theorem 17 proves that *every NSM under exchangeability must be a nonincreasing process and thus powerless*. The way out is to look beyond NSMs, at an object called a safe e-process, which we introduce next. Safe e-processes are intimately related to NSMs — all NSMs are safe e-processes, and all admissible safe e-processes are infima of NSMs [22] — but the two concepts are different, and our paper seeks to explore this gap in detail and exploit it successfully.

1.1 Safe e-processes and anytime-valid p-processes

We define a $\overline{\mathcal{Q}}$ -safe e-process to be a nonnegative sequence of adapted random variables $(E_t)_{t \geq 0}$ such that

$$\sup_{\mathbf{Q} \in \overline{\mathcal{Q}}} \sup_{\tau} \mathbf{E}_{\mathbf{Q}}[E_\tau] \leq 1, \tag{2}$$

where the second supremum is over all stopping times τ , possibly infinite. Above, we interpret $E_\infty := \limsup_{t \rightarrow \infty} E_t$ to accommodate potentially infinite stopping times. (As mentioned earlier, the filtration (\mathcal{F}_t) defined previously is fixed and implicit.) Howard et al. [12] do not use the term e-process, but introduce study exponential processes that are upper bounded by composite nonnegative supermartingales, which is a characterizing property of e-processes. Related notions have been studied by Shafer et al. [31], Grünwald et al. [10], Ramdas et al. [22], Vovk and Wang [36], amongst others.

Remark 1 (on terminology). *An ‘e-variable’ is a nonnegative random variable that has expectation at most one under the null [36]; in other words E_t is an e-variable for every t , while we call (E_t) as an e-process due to the additional property (2) that insists that stopped e-processes are also e-variables, meaning that E_τ is an e-variable for any stopping time τ . The value realized by an e-variable will be called an e-value. While recent works in this rapidly evolving area have sometimes blurred the distinctions between e-processes, e-variables and e-values, we believe that our choices of terminology may be more suitable to emphasize sequential problem setups involving stopping times.*

Large e-values encode evidence against the null, and it is easy to check that the stopping time

$$\kappa_\alpha := \inf \left\{ t \geq 1 : E_t \geq \frac{1}{\alpha} \right\} \tag{3}$$

results in a level α sequential test by applying Markov's inequality to the stopped e-process. More details can be found in [22], who also show that the sequence (\mathbf{p}_t) defined by $\mathbf{p}_t := \inf_{s \leq t} 1/E_s$ is an anytime-valid (or $\overline{\mathcal{Q}}$ -valid) p-process, meaning that

$$\sup_{\mathbf{Q} \in \overline{\mathcal{Q}}} \sup_{\tau} \mathbf{Q}(\mathbf{p}_{\tau} \leq \alpha) \leq \alpha, \quad \text{for all } \alpha \in [0, 1]. \quad (4)$$

Here the second supremum is taken over all stopping times τ . Following the above remark, a p-variable is a random variable whose distribution is stochastically larger than uniform under the null. Thus, a stopped p-process is a p-variable, and we call its realization as a p-value (a number between 0 and 1).

Since such $\overline{\mathcal{Q}}$ -safe e-processes yield both sequential tests and anytime-valid p-processes, we focus on constructing e-processes for the rest of this paper. As a matter of convention, we always use κ_{α} to denote the above stopping time, that is the one that thresholds a safe e-process at level $1/\alpha$, while τ denotes a generic stopping time.

Power one and consistency. We call an e-process powerful if the corresponding test is powerful, and of course we desire a test that is consistent, meaning that its power goes to one with the sample size. Formally, a level α sequential test τ_{α} has asymptotically power one against some family $\mathcal{P} \supset \overline{\mathcal{Q}}$ if

$$\inf_{\mathbf{P} \in \mathcal{P} \setminus \overline{\mathcal{Q}}} \mathbf{P}(\tau_{\alpha} < \infty) = 1.$$

(Henceforth, if we mention some alternative \mathcal{P} , it is understood that we desire power against $\mathcal{P} \setminus \overline{\mathcal{Q}}$.)

A $\overline{\mathcal{Q}}$ -safe e-process (E_t) is said to be consistent, or power one, if its associated sequential test κ_{α} from (3) is asymptotically power one at any level, meaning that

$$\text{for all } \alpha \in (0, 1), \quad \inf_{\mathbf{P} \in \mathcal{P} \setminus \overline{\mathcal{Q}}} \mathbf{P}(\kappa_{\alpha} < \infty) = 1, \quad (5a)$$

or equivalently,

$$\limsup_{t \rightarrow \infty} E_t = \infty, \quad \mathbf{P}\text{-almost surely, for every } \mathbf{P} \in \mathcal{P} \setminus \overline{\mathcal{Q}}. \quad (5b)$$

1.2 Convex hulls and de Finetti's theorem

Let the null set \mathcal{Q} consist of all product distributions $\text{Ber}(p)^{\infty}$ for some $p \in [0, 1]$. Note that \mathcal{Q} is a rich composite class of parametric distributions, whose convex hull is $\overline{\mathcal{Q}}$, which is a result well known as de Finetti's theorem [17].

Any sequential test for \mathcal{Q} is also valid for $\overline{\mathcal{Q}}$, a fact that we record below, proved in [22].

Proposition 2. *The properties of type-1 error control (1) and safety are closed under the convex hull, meaning that any $\overline{\mathcal{Q}}$ -safe e-process is also $\overline{\mathcal{Q}}$ -safe, and any level α sequential test for \mathcal{Q} is also valid for $\overline{\mathcal{Q}}$.*

As a consequence, we may restrict our attention to developing a $\overline{\mathcal{Q}}$ -safe e-process, and invoke the above fact to step from the i.i.d. setting to the exchangeable setting. This will be our approach in the rest of this paper. Another consequence is that testing the null \mathcal{Q} against the alternative $\overline{\mathcal{Q}}$ is futile; safe and consistent e-processes do not exist and neither do valid, power-one tests.

Remark 3. *To avoid confusion, we note that the convex combination of $\mathbf{Q}, \mathbf{Q}' \in \mathcal{Q}$ must be carefully interpreted. For example, if $\mathbf{Q} = \text{Ber}(0.3)^{\infty}$ and $\mathbf{Q}' = \text{Ber}(0.7)^{\infty}$ then a draw from $(\mathbf{Q} + \mathbf{Q}')/2$ produces either a sequence with 70% zeros or a sequence with 70% ones, each with probability half, while it produces a sequence with equal number of zeros and ones with probability zero. Contrast this with the fact that a draw from $((\text{Ber}(0.3) + \text{Ber}(0.7))/2)^{\infty}$ is equally likely to produce a zero or a one at every point in time, and thus produces a sequence with equal zeros and ones with probability one. Thus, one must take care to differentiate between $((\text{Ber}(0.3) + \text{Ber}(0.7))/2)^{\infty}$, which is not in the convex hull of \mathbf{Q} and \mathbf{Q}' , and $(\mathbf{Q} + \mathbf{Q}')/2$, which is. Later, we will see that the former lies in the closed fork-convex hull of \mathbf{Q} and \mathbf{Q}' .*

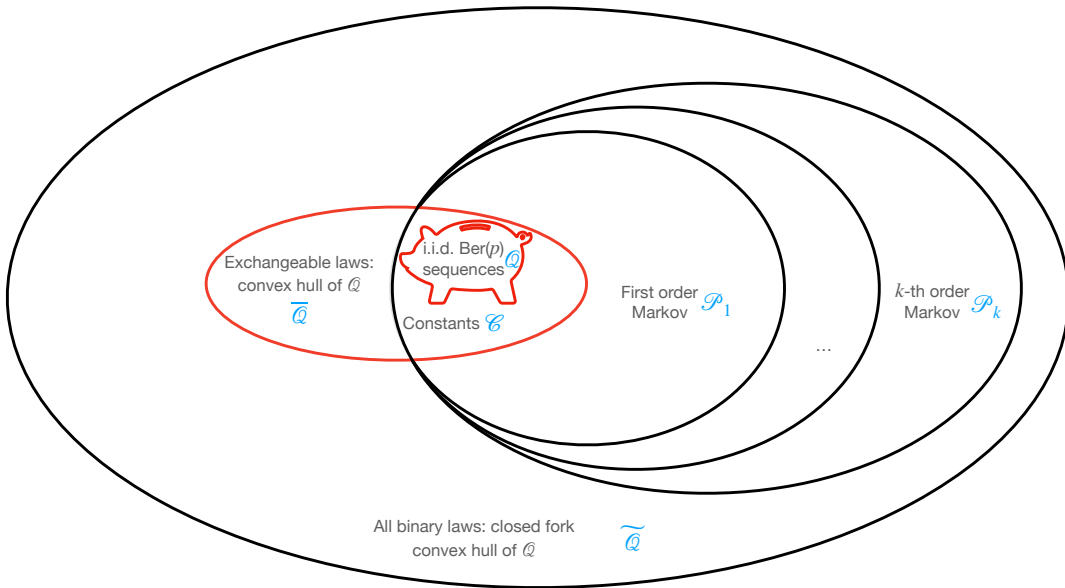


Figure 1: Various classes of distributions over infinite binary sequences encountered in this paper. See Proposition 4 and Theorem 17 for the proof of the displayed set inclusions.

It is impossible to have a powerful test for $\overline{\mathcal{Q}}$ against its complement $\overline{\mathcal{Q}}^c$, since the alternative is too rich and consists of too many distributions that are too close to $\overline{\mathcal{Q}}$, meaning that there are too many ways to violate exchangeability. For example, it should be apparent to the reader that if the first coin has bias p_1 and every other coin has bias $p \neq p_1$, then the resulting sequence is not exchangeable but we would never be able to reliably detect this deviation. This example relies on ensuring that the information required to detect a deviation from the null is exhausted early on in the sequence. To avoid such pathologies it is necessary to restrict the alternative class in some meaningful way. Markovian alternatives are an attractive choice, balancing the needs of relevant practical motivation, tractable mathematical structure, succinct probabilistic description, and intuitive aesthetic appeal.

1.3 Time-homogeneous Markovian alternatives

We focus on the setting of first-order Markov alternatives \mathcal{P}_1 , and return to address higher-order (and even more general) alternatives later. To describe \mathcal{P}_1 more formally, each $P \in \mathcal{P}_1$ represents a time-homogeneous first-order Markov process with parameters $p_{1|0}$ and $p_{1|1}$. Here we abbreviate $p_{0|0} = 1 - p_{1|0}$ and $p_{0|1} = 1 - p_{1|1}$. For arbitrary $k \in \mathbb{N}$, we will also consider \mathcal{P}_k , the set of k -th order Markov processes.

Let \mathcal{C} represent the distributions that result in constant sequences (X_t) , i.e., where $X_1 = X_2 = \dots$. These include, of course, the distributions of deterministic sequences 0^∞ and 1^∞ . The class \mathcal{C} also includes their mixtures, which are distributions that first set $X_1 \sim \text{Ber}(p)$ for some p and then set $X_t = X_1$ for all $t > 1$. In other words, conditional on X_1 , the remaining observations are i.i.d. $X_t \sim \text{Ber}(X_1)$.

Proposition 4. *For any $\{0, 1\}$ -valued stochastic process (X_t) , the following two conditions are equivalent:*

- (i) (X_t) is k -th order Markov (with respect to its natural filtration) for some $k \in \mathbb{N}$ and exchangeable.
- (ii) (X_t) is either constant (as a function of time) or i.i.d.

Said differently, $\overline{\mathcal{Q}} \cap \mathcal{P}_k = \mathcal{C} \cup \mathcal{C}$ for any $k \in \mathbb{N}$.

This relation and more are illustrated in Figure 1. The proof and some additional remarks can be found in Appendix B.

1.4 Wald’s sequential likelihood (or probability) ratio test

Let the likelihood under a particular $\mathbb{Q} \in \mathcal{Q}$, where $\mathbb{Q} = \text{Ber}(p)^\infty$, be represented by

$$Q_t \equiv Q_p(X_1, \dots, X_t) := (1 - p)^{n_0} p^{n_1},$$

where $n_0 = n_0(t)$ and $n_1 = n_1(t)$ represent the number of zeros and ones seen up to time t . The likelihood associated to $\mathbb{P} \in \mathcal{P}_1$, say, for simplicity with $\mathbb{P}(X_1 = 1) = 1/2$, is given by

$$P_t \equiv P_{p_{1|0}, p_{1|1}}(X_1, \dots, X_t) := \frac{1}{2} \prod_{s=2}^t p_{X_s|X_{s-1}} = \frac{1}{2} p_{0|0}^{n_{0|0}} p_{1|0}^{n_{1|0}} p_{0|1}^{n_{0|1}} p_{1|1}^{n_{1|1}}.$$

Here $n_{1|0} \equiv n_{1|0}(t)$ denotes the total number of ones following zeros up to time t , etc. (We omit the dependence on t to ease notational load.)

For a point null $\mathbb{Q} \in \mathcal{Q}$ and point alternative $\mathbb{P} \in \mathcal{P}_1$ with $\mathbb{P} \neq \mathbb{Q}$, Wald’s sequential likelihood (or probability) ratio test (SLRT or SPRT) [39] yields a power-one test. The likelihood ratio process, i.e., (P_t/Q_t) , is a \mathbb{Q} -martingale starting at one and thus a \mathbb{Q} -safe e-process, and the resulting sequential test $\inf\{t \geq 0 : P_t/Q_t \geq 1/\alpha\}$ is a level α test under \mathbb{Q} that is known — up to a discrete-time issue called “boundary overshoot” — to have the smallest expected stopping time under \mathbb{P} [41].

For composite nulls and alternatives, the SLRT cannot be directly applied. The *mixture SLRT* integrates over the alternatives using a “prior” distribution (or, more appropriately, mixture distribution, to avoid any Bayesian interpretations of our frequentist statements). However, this only works for composite alternatives, since mixing over the null set does not yield a safe e-process or the desired type-1 error control property in (1). (We note here that, interestingly, the GROW e-variables of [10] are ratios of mixtures, though they are safe only for a fixed sample size.) The generalized SLRT maximizes the likelihood under both null and alternative, but this also does not yield a martingale or a safe e-process. In both cases, it is difficult to find a threshold for the resulting process that achieves type-1 error control in (1), since the SLRT’s choice of $1/\alpha$ does not suffice. For a discussion of techniques and results for this approach we refer to [16].

1.5 Nonnegative supermartingales (NSMs) and an impossibility result

Despite the above apparent difficulties in generalizing the SLRT to yield an e-process, it has been recently established that nonnegative (super)martingales play a fundamental role in the design of admissible sequential tests (and the construction of admissible safe e-processes), even for composite nonparametric nulls [12, 22]. In anticipation of the results to follow, it is useful to set up some relevant notation. In what follows, a process $(M_t)_{t \geq 0}$ will be called a \mathbb{Q} -NM if it is a nonnegative martingale with initial value one, that is, (M_t) is adapted to (\mathcal{F}_t) , $M_0 = 1$ and $\mathbb{E}_{\mathbb{Q}}[M_t | \mathcal{F}_s] = M_s \geq 0$ for any $s \leq t$. Such processes are called test martingales by Shafer et al. [31]. If (M_t) is a \mathbb{Q} -NM simultaneously for every $\mathbb{Q} \in \mathcal{Q}$, then we will call it a \mathcal{Q} -NM. If the equality above is replaced by an inequality \leq , then we will call it a \mathbb{Q} -NSM or \mathcal{Q} -NSM (nonnegative supermartingale).

An appropriate variant of the optional stopping theorem [7] implies that for any \mathcal{Q} -NSM (M_t) and any stopping time τ (potentially infinite), the stopped process has expectation at most one, or in other words

$$\sup_{\mathbb{Q} \in \mathcal{Q}} \sup_{\tau} \mathbb{E}_{\mathbb{Q}}[M_\tau] \leq 1,$$

where the second supremum is over *all* stopping times τ , potentially infinite [22, Section 3]. Indeed, for any \mathcal{Q} -NSM, $M_\infty := \lim_{t \rightarrow \infty} M_t$ is a well defined random variable, and $\sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[M_\infty] \leq 1$. The correspondence with the definition of a safe e-process (recall (2)) is not coincidental — to construct a \mathcal{Q} -safe e-process, it suffices to construct a \mathcal{Q} -NSM. However, we claim the following.

Proposition 5. *Every \mathcal{Q} -NSM is also a \mathcal{P}_1 -NSM (recall that \mathcal{Q} and \mathcal{P}_1 contain all i.i.d. and first-order Markov distributions respectively). In fact, every \mathcal{Q} -NSM is also a \mathcal{P}_k -NSM for every $k \in \mathbb{N}$, and is also a \mathcal{P} -NSM for any set \mathcal{P} . In other words, any \mathcal{Q} -safe e-process with nontrivial power cannot be a \mathcal{Q} -NSM, since the latter is powerless against any alternative \mathcal{P} by virtue of being a \mathcal{P} -NSM.*

This proposition will follow from Theorem 17. This paper is as much about understanding the above negative result, as about providing a positive result. In other words, this result probes at the “gap” between a \mathcal{Q} -safe e-process and a \mathcal{Q} -NSM. The former is a much weaker property than the latter. While the latter suffices for the former, it is by no means necessary, as recently observed in a more abstract setup [22]. Indeed, a \mathcal{Q} -NSM (N_t) satisfies the much stronger “conditional” property that

$$\mathbb{E}_{\mathbb{Q}}[N_{\tau}|\mathcal{F}_s] \leq N_{\tau \wedge s} \text{ for every } \mathbb{Q} \in \mathcal{Q} \text{ and stopping time } \tau,$$

which implies the earlier mentioned optional stopping result. In fact, the above property is satisfied if and only if (N_t) is a \mathcal{Q} -NSM, but the e-process property (2) can be satisfied even by processes that are upper bounded by \mathcal{Q} -NSMs for each $\mathbb{Q} \in \mathcal{Q}$, but are not themselves \mathcal{Q} -NSMs. It is exactly this gap that we will exploit going forward.

We provide a geometrical characterization of the above phenomenon: essentially, we will show that the above proposition is true because \mathcal{P}_k lies within the “fork-convex hull” of \mathcal{Q} , and we prove that this hull preserves the NSM property of a process. In fact, Theorem 17 proves that the fork-convex hull of \mathcal{Q} is all-encompassing, containing every law over binary sequences. Thus, a \mathcal{Q} -NSM yields a powerless test against \mathcal{P}_k since it is automatically and unintentionally safe under the alternative as well as under the null. Along the way, we will encounter other friends from martingale theory, such as the Snell envelope, which previously has not played a prominent role in the mathematical treatment of sequential testing.

1.6 A powerful \mathcal{Q} -safe e-process that is not a \mathcal{Q} -NSM

Ramdas et al. [22] show that any \mathcal{Q} -safe e-process is dominated by a \mathcal{Q} -safe e-process (E_t) of the form

$$E_t := \inf_{\mathbb{Q} \in \mathcal{Q}} M_t^{\mathbb{Q}},$$

where $(M_t^{\mathbb{Q}})$ is a \mathbb{Q} -NM. As mentioned before, each limiting variable $M_{\infty}^{\mathbb{Q}}$ is well defined and has expectation at most one, and thus, $E_{\infty} := \limsup_{t \rightarrow \infty} E_t$ also has expectation at most one. The safety property immediately holds at infinite times as well, meaning that $\sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[E_{\infty}] \leq 1$ and $\sup_{\mathbb{Q} \in \mathcal{Q}} \sup_{\tau} \mathbb{E}_{\mathbb{Q}}[E_{\tau}] \leq 1$, where the second supremum is over all stopping times τ . To avoid too much suspense before we get into these subtle new concepts, we first present our solution immediately, and delay its derivation to the next section. To this end, define

$$R_t := \frac{\Gamma(n_{0|0} + \frac{1}{2}) \Gamma(n_{0|1} + \frac{1}{2}) \Gamma(n_{1|0} + \frac{1}{2}) \Gamma(n_{1|1} + \frac{1}{2})}{2\Gamma(\frac{1}{2})^4 \Gamma(n_{0|0} + n_{1|0} + 1) \Gamma(n_{0|1} + n_{1|1} + 1)} / \left(\left(\frac{n_1}{t}\right)^{n_1} \left(\frac{n_0}{t}\right)^{n_0} \right), \quad (6)$$

where Γ denotes the usual gamma function.

Theorem 6. *The process (R_t) is a \mathcal{Q} -safe e-process, and thus thresholding it at level $1/\alpha$ yields a level α sequential test κ_{α} for exchangeability.*

The first three subsections of Section 2 is devoted to the justification of the above theorem (its proof is implicit in its derivation). Later, Theorem 7 proves that this test has power one for first-order Markov alternatives, i.e., (5a) holds, and quantifies the rate of growth. Above, (R_t) is not itself a \mathcal{Q} -NSM, but is nevertheless upper bounded by a (different!) \mathbb{Q} -NM for every $\mathbb{Q} \in \mathcal{Q}$, resulting in it being a \mathcal{Q} -safe e-process. This idea is enabled by bringing together the method of mixtures (using Jeffreys’ prior) for combining the composite alternative, with the maximum likelihood under the composite null. Beyond showing that it has power one, one can quantify that it has rate-optimal power by utilizing a regret bound from universal coding. The next section, among others, proves the above theorem, after which we turn to defining fork-convexity. Then we provide some numerical simulations and conclude with a discussion about this paper’s approach compared to other possible approaches to the problem.

1.7 A summary of this paper’s contributions

This paper proposes a nontrivial and powerful test for exchangeability, a fundamental and easy-to-state problem that exposes rich geometric and probabilistic structure. The most directly related work is that of Vovk [35], who proposed a very different approach based on conformal prediction, discussed in Section 5.3.

The problem of testing exchangeability also serves as a test-bed to carefully examine the gap between processes that are composite NSMs, and processes that are composite safe e-processes. The former is a special case of the latter and the latter is a necessary ingredient for the former [22], but the difference between these concepts has remained somewhat implicit in past work [12]. This work constitutes the first detailed investigation of a problem where nontrivial NSMs do not exist (Proposition 5, Theorem 17) but powerful safe e-processes do exist (Theorem 6, Theorem 7).

In the process, we prove many results that are of independent interest. For example, the NSM property is closed under “fork-convex combinations” of distributions (Lemma 12, Theorem 13). We prove that a composite safe e-process can be improved upon by a composite NSM only if the set of distributions is fork-convex, and this improvement is given by a composite Snell envelope (Theorem 14, Corollary 16).

Section 4 examines the empirical behavior of our safe e-process in a simulation suite that includes both basic sanity checks and rather subtle and nontrivial alternatives. Finally, Section 5.4 suggests a way to interpret our e-process in the language of game-theoretic protocols of Shafer and Vovk [30].

2 Jeffreys’ mixture meets maximum likelihood

As briefly mentioned earlier, if the null set was a singleton, say corresponding to $\text{Ber}(p)^\infty$, and the alternative was also a singleton, such as when (X_t) deterministically alternates between 0 and 1, then Wald’s sequential likelihood ratio test [39] would immediately yield a solution to the problem at hand. To elaborate, let f_p denote the probability mass function of a $\text{Ber}(p)$ random variable and let L_{01} denote the likelihood function under the alternative:

$$L_{01}(X_1, \dots, X_t) := \begin{cases} 1 & \text{if } (X_1, X_2, \dots) = (0, 1, 0, 1, \dots); \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any point null (indexed by p) define the following likelihood ratio:

$$R_t^p := \frac{L_{01}(X_1, \dots, X_t)}{\prod_{s=1}^t f_p(X_s)}, \quad \text{which equals} \quad \begin{cases} \frac{1}{p^{\lfloor t/2 \rfloor} (1-p)^{\lceil t/2 \rceil}} & \text{with } \text{Ber}(p)^\infty\text{-probability } p^{\lfloor t/2 \rfloor} (1-p)^{\lceil t/2 \rceil}; \\ 0 & \text{with } \text{Ber}(p)^\infty\text{-probability } 1 - p^{\lfloor t/2 \rfloor} (1-p)^{\lceil t/2 \rceil}. \end{cases}$$

It is easy to check that (R_t^p) is a $\text{Ber}(p)^\infty$ -NM, and thus a $\text{Ber}(p)^\infty$ -safe e-process, and κ_α from (3) yields a valid level α sequential test that coincides with Wald’s original proposal [39]. The question is how to generalize this approach to deal with a composite null and a composite alternative in a computationally tractable and statistically powerful manner.

The following observation deals the first blow: the only process that is a nonnegative martingale under every i.i.d. Bernoulli sequence is one that is constant. In other words, the only \mathcal{Q} -NM is such that $M_t = 1$ for all $t \in \mathbb{N}_0$. This obviously results in a powerless test. So we then turn our attention to constructing a \mathcal{Q} -NSM, or a test supermartingale. Unfortunately this approach is dealt a fatal blow by Proposition 5. As alluded to in the introduction, we cannot employ mixtures in both numerator and denominator because it violates safety by lowering the denominator too much, and we cannot maximize the likelihood in the numerator and denominator because it violates safety by raising the numerator too much.

Our proposal combines a suitably chosen mixture in the numerator with maximum likelihood in the denominator, thus avoiding both pitfalls.

2.1 Dealing with the composite null via maximum likelihood

To start, let us return to the point alternative described above (alternating 0 and 1), and just handle the composite null using maximum likelihood estimation, as proposed in universal inference [42]. To elaborate,

observe that

$$R_t^{\text{ML}} := \inf_{p \in [0,1]} R_t^p = \frac{\text{likelihood under the point alternative}}{\text{maximum likelihood under the null}}$$

is a \mathcal{Q} -safe e -value. Indeed, suppose the data truly comes from $\text{Ber}(p^*)^\infty$ for an unknown p^* . Then, it is obvious that $R_t^{\text{ML}} \leq R_t^{p^*}$, where the latter process is a $\text{Ber}(p^*)^\infty$ -NM. Thus, for any $\mathcal{Q} \in \mathcal{Q}$ (corresponding to some $p^* \in [0, 1]$) and any stopping time τ , we have

$$\mathbb{E}_{\mathcal{Q}}[R_\tau^{\text{ML}}] \leq \mathbb{E}_{\mathcal{Q}}[R_\tau^{p^*}] \leq 1,$$

where the last step invokes the optional stopping theorem for the $\text{Ber}(p^*)^\infty$ -NM ($R_t^{p^*}$).

To see that the resulting test has good power, note that under the alternative, $p = 1/2$ uniquely achieves the above infimum at any even time $t \in 2\mathbb{N}$, in which case the denominator equals $(1/2)^t$ and the numerator equals one. Thus, $R_t^{\text{ML}} = 2^t$ at even times, and the test $\inf\{t \geq 0 : R_t^{\text{ML}} \geq 1/\alpha\}$ is a valid level α sequential test that stops either at time $\lceil \ln(1/\alpha)/\ln(2) \rceil$ or at time $\lceil \ln(1/\alpha)/\ln(2) \rceil + 1$.

To recap, despite the fact that we cannot find a \mathcal{Q} -NM, (R_t^{ML}) is a powerful \mathcal{Q} -safe e -process against the considered point alternative. This test takes the ratio of the likelihood under the alternative to the maximum likelihood under the null. Next, we detail how to handle the composite alternative \mathcal{P} when testing a point null in \mathcal{Q} .

2.2 Dealing with the composite alternative using a Jeffreys' mixture

A standard way of dealing with composite alternatives is to “mix” over them by choosing an appropriate mixture distribution. This is sometimes called the Laplace method, pseudo-maximization, or Robbins' method of mixtures [25, 23, 27, 24, 26]. It was first utilized successfully primarily in parametric settings, but has recently re-emerged as a powerful tool for nonparametric inference [5, 13, 16, 43]. This mixture is sometimes called a “prior”, but it must be made clear that this is not a prior in the Bayesian sense — any other distribution also suffices to serve as a mixture in terms of maintaining safety, and no extra assumption is made on the true data when such a “working prior” is employed. The choice of mixture does not affect safety, but it does affect computability and power. In the following, we work with a very particular choice of mixture because it yields closed-form expressions, and it yields small constants in regret bounds from the universal coding literature (which has implications for power).

Taking independent Jeffreys' priors (with densities $w(\theta) := 1/(\pi\sqrt{\theta(1-\theta)})$) for $p_{1|0}$ and $p_{1|1}$, we obtain the mixture likelihood

$$\begin{aligned} \mathbb{P}_{w \times w}(X_1, \dots, X_t) &:= \int \mathbb{P}_{p_{1|0}, p_{1|1}}(X_1, \dots, X_t) w(p_{1|0}) w(p_{1|1}) d(p_{1|0}, p_{1|1}) \\ &= \frac{\Gamma(n_{0|0} + \frac{1}{2}) \Gamma(n_{0|1} + \frac{1}{2}) \Gamma(n_{1|0} + \frac{1}{2}) \Gamma(n_{1|1} + \frac{1}{2})}{2\Gamma(\frac{1}{2})^4 \Gamma(n_{0|0} + n_{1|0} + 1) \Gamma(n_{0|1} + n_{1|1} + 1)}. \end{aligned}$$

Here we have taken the mixture over all first-order Markov distributions $\mathbb{P}_{p_{1|0}, p_{1|1}}$, under which $\{X_1 = 1\}$ has probability $1/2$.

Thus, for any point null represented by $\text{Ber}(p)^\infty$, we can define the mixture likelihood ratio

$$R_t^{\text{JP}, p} := \frac{\mathbb{P}_{w \times w}(X_1, \dots, X_t)}{\prod_{s=1}^t f_p(X_s)} = \frac{\text{Jeffreys' mixture over the alternative}}{\text{likelihood under the point null}}. \quad (7)$$

Using Fubini's theorem to swap integrals, it is easy to check that $R_t^{\text{JP}, p}$ is a $\text{Ber}(p)^\infty$ -NM, and the corresponding sequential test is Wald's usual mixture SLRT [40]. Note that it is the very particular form of this mixture that yields a closed form expression and thus a computationally feasible test. However, we do not use this mixture just for computational reasons; as we detail soon, combining it with the earlier maximum likelihood idea also yields a statistically near-optimal power.

2.3 Combining Jeffreys' mixture with maximum likelihood

Using, as in the previous example, that the likelihood under the null is maximised at $p = n_1/t$, where it evaluates to $(n_1/t)^{n_1} (n_0/t)^{n_0}$, we find that

$$R_t := \frac{\text{Jeffreys' mixture over the alternative}}{\text{maximum likelihood under the null}}$$

reduces to the expression in (6). This is a \mathcal{Q} -safe e -value by combining the arguments used for the safety of $(R_t^{\text{JP},p})$ and (R_t^{ML}) : swapping the maximum likelihood with the (unknown) true likelihood, and then employing Fubini's theorem.

We remark that any prior above would have yielded a \mathcal{Q} -safe e -value, and in fact any Beta prior would have yielded one in closed form, but the Jeffreys' mixture above allows us to invoke an appropriate optimal regret bound from the universal coding literature [20] to Markov sources (see [33] for a discussion of the resulting optimality):

$$\begin{aligned} R_t &\geq \frac{\frac{1}{2} \left(\frac{n_{1|0}}{n_{1|0} + n_{0|0}} \right)^{n_{1|0}} \left(\frac{n_{0|0}}{n_{1|0} + n_{0|0}} \right)^{n_{0|0}} \left(\frac{n_{1|1}}{n_{1|1} + n_{0|1}} \right)^{n_{1|1}} \left(\frac{n_{0|1}}{n_{1|1} + n_{0|1}} \right)^{n_{0|1}}}{\left(\frac{n_1}{t} \right)^{n_1} \left(\frac{n_0}{t} \right)^{n_0}} \times e^{-\frac{1}{2} \ln(n_{1|0} + n_{0|0}) - \frac{1}{2} \ln(n_{1|1} + n_{0|1}) - O(1)} \\ &= \frac{\text{maximum likelihood of Markov model}}{\text{maximum likelihood of Bernoulli model}} \times e^{-\ln t - O(1)}. \end{aligned} \quad (8)$$

That is, (R_t) starts gathering evidence against the null if the maximum likelihood for the first-order Markov chain outperforms the maximum likelihood for the Bernoulli model by a factor of order t . Note that this is a small hurdle to overcome, as the first term is growing exponentially fast in t when the data are explained better by a Markov model, as argued next.

2.4 A pathwise theorem on the power of (R_t)

Although our main focus is on first-order Markov alternatives, (R_t) actually has power under much more general alternatives. Consider any binary sequence for which the following limits exist:

$$\alpha := \lim_{t \rightarrow \infty} \frac{n_{1|1}}{t}, \quad \beta := \lim_{t \rightarrow \infty} \frac{n_{0|0}}{t}, \quad (9a)$$

$$\gamma := \lim_{t \rightarrow \infty} \frac{n_{1|0}}{t} = \lim_{t \rightarrow \infty} \frac{n_{0|1}}{t}, \quad p := \lim_{t \rightarrow \infty} \frac{n_1}{t} = \lim_{t \rightarrow \infty} \frac{n_{1|1} + n_{0|1}}{t}. \quad (9b)$$

The theorem that follows is a ‘‘pathwise’’ result, holding for any sequence where the above limits exist and satisfy an additional weak condition, described later. These limits exist, for example, under *any* k -th order Markov alternative (but may be random unless suitable irreducibility properties hold). Note that the equality of the two expressions for γ follows from the fact that $n_{1|0}$ and $n_{0|1}$ can differ by at most one. Similarly, n_1 differs from $n_{1|1} + n_{0|1}$ by at most one, which explains why the two expressions for p are equal. Moreover, note that one has the relations

$$p = \alpha + \gamma = 1 - \beta - \gamma = \frac{p_{1|0}}{p_{1|0} + p_{0|1}},$$

where we define

$$p_{1|1} := \frac{\alpha}{p}, \quad p_{0|1} := \frac{\gamma}{p}, \quad p_{1|0} := \frac{\gamma}{1-p}, \quad p_{0|0} := \frac{\beta}{1-p}.$$

Here we use the convention $0/0 := 1$. In the first-order Markov case these parameters are simply the transition probabilities.

Within the class of first-order Markov chains, the special case of i.i.d. Bernoulli data is characterized by the restriction $p_{1|1} = p_{1|0}$ (and hence $p_{0|1} = p_{0|0}$), or equivalently

$$\alpha(1-p) = \gamma p \quad (10)$$

(and hence $\gamma(1-p) = \beta p$). For more general models, such as higher-order Markov chains, the restriction $\alpha(1-p) = \gamma p$ no longer characterizes the i.i.d. Bernoulli case. This is illustrated in Example 9, where it is also shown that (R_t) can be powerless against such alternatives (that is, under higher-order Markov alternatives that appear to be Bernoulli when summarized by only first order transition parameters). Nonetheless, as explained in Remark 8 below, in the suitably ergodic non-Bernoulli k -th order Markov case one still has $\alpha(1-p) \neq \gamma p$ generically, and therefore achieves power against such alternatives thanks to the following theorem.

Theorem 7. *For any data sequence, if the limits α, β, γ, p in (9) exist, then $\lim_{t \rightarrow \infty} \ln R_t / t = r^*$, where*

$$\begin{aligned} r^* &:= p \left(\ln \frac{1}{p} - p_{0|1} \ln \frac{1}{p_{1|0}} - p_{1|1} \ln \frac{1}{p_{1|1}} \right) + (1-p) \left(\ln \frac{1}{1-p} - p_{0|0} \ln \frac{1}{p_{0|0}} - p_{1|0} \ln \frac{1}{p_{0|1}} \right) \\ &= \gamma \ln \frac{\gamma}{1-p} + \beta \ln \frac{\beta}{1-p} + \alpha \ln \frac{\alpha}{p} + \gamma \ln \frac{\gamma}{p} - p \ln p - (1-p) \ln(1-p). \end{aligned} \quad (11)$$

Note that the second expression is a difference in entropies between first-order Markovian and Bernoulli sources with the corresponding parameters. Furthermore, if $\alpha(1-p) \neq \gamma p$ we have $r^* > 0$, that is, (R_t) increases to infinity exponentially fast. In fact, $r^* = 0$ if and only if $\alpha(1-p) = \gamma p$. Finally, whenever the exchangeability null is true, $\lim_{t \rightarrow \infty} R_t = 0$ almost surely, in addition to (R_t) being $\overline{\mathcal{Q}}$ -safe.

Note that whenever the null is true, we necessarily have $r^* = 0$. However, $r^* = 0$ does not by itself characterize power or consistency; it only characterizes when (R_t) grows exponentially. It is possible to construct certain “borderline” examples where $r^* = 0$, but (R_t) still increases to infinity (our test is powerful and consistent), albeit at a subexponential rate. We describe and explore such an example in the numerical simulations in Section 4.3.

In the case of a non-Bernoulli first-order Markov alternative, the condition of the theorem is satisfied for almost every realization of the data (that is, with probability one), provided the chain does not have any absorbing states. Indeed, with no absorbing states the Markov chain is recurrent and hence the ergodic theorem applies. This condition is necessary for a power one test, because if (say) 1 is an absorbing state then there is positive probability of seeing only ones (unless, of course, the Markov chain starts at 0 with probability one). This is indistinguishable from a realization of an i.i.d. $\text{Ber}(1)$ sequence.

Remark 8. *Any k -th order Markov chain is characterized by its transition probabilities $p_{1|\mathbf{s}} \in [0, 1]$ where \mathbf{s} ranges over all prefixes of length k . Note that this determines $p_{0|\mathbf{s}} = 1 - p_{1|\mathbf{s}}$, which we thus do not count as separate parameters. The quantities α, β, γ, p defined previously, when they are deterministic, are real analytic functions of the transition probabilities. This implies that the condition $\alpha(1-p) \neq \gamma p$ of Theorem 7 holds “generically” (in the topological sense, i.e., for an open dense set of parameters). Indeed, a real analytic function is either identically zero, or nonzero on an open dense set, and we know that $\alpha(1-p) = \gamma p$ does not hold identically (i.e., for all choices of parameters $p_{1|\mathbf{s}}$). This confirms that $\alpha(1-p) \neq \gamma p$ must hold generically.*

Proof of Theorem 7. Let $\ell(t)$ denote the logarithm of the ratio between the maximum likelihood for the first-order Markov model and the maximum likelihood for the Bernoulli model. We then have

$$\begin{aligned} \frac{\ell(t)}{t} &= \frac{n_{1|0}}{t} \ln \frac{n_{1|0}}{n_{1|0} + n_{0|0}} + \frac{n_{0|0}}{t} \ln \frac{n_{0|0}}{n_{1|0} + n_{0|0}} + \frac{n_{1|1}}{t} \ln \frac{n_{1|1}}{n_{1|1} + n_{0|1}} \\ &\quad + \frac{n_{0|1}}{t} \ln \frac{n_{0|1}}{n_{1|1} + n_{0|1}} - \frac{n_1}{t} \ln \frac{n_1}{t} - \frac{n_0}{t} \ln \frac{n_0}{t}. \end{aligned}$$

Sending $t \rightarrow \infty$ the right-hand side converges to

$$\gamma \ln \frac{\gamma}{1-p} + \beta \ln \frac{\beta}{1-p} + \alpha \ln \frac{\alpha}{p} + \gamma \ln \frac{\gamma}{p} - p \ln p - (1-p) \ln(1-p),$$

which is precisely (11). Since $\ell(t) - \ln t + O(1) \leq \ln R_t \leq \ell(t)$ by (8) we deduce that $t^{-1} \ln R_t$ converges to r^* as claimed. Re-arranging (11) we get

$$r^* = p \left(\ln \frac{1}{p} - \frac{\gamma}{p} \ln \frac{1-p}{\gamma} - \frac{\alpha}{p} \ln \frac{p}{\alpha} \right) + (1-p) \left(\ln \frac{1}{1-p} - \frac{\beta}{1-p} \ln \frac{1-p}{\beta} - \frac{\gamma}{1-p} \ln \frac{p}{\gamma} \right). \quad (12)$$

If $\alpha(1-p) = \gamma p$ then one easily verifies that $r^* = 0$. Next, suppose $\alpha(1-p) \neq \gamma p$. Since $\alpha + \gamma = p$, we may bound the first parenthesized expression in (12) using Jensen's inequality,

$$\ln \frac{1}{p} - \frac{\gamma}{p} \ln \frac{1-p}{\gamma} - \frac{\alpha}{p} \ln \frac{p}{\alpha} > \ln \frac{1}{p} - \ln \left(\frac{\gamma}{p} \frac{1-p}{\gamma} + \frac{\alpha}{p} \frac{p}{\alpha} \right) = 0,$$

where the inequality is strict because $(1-p)/\gamma \neq p/\alpha$. By the same token, the second parenthesized expression in (12) is strictly positive as well. In particular, $r^* > 0$. Finally, if the exchangeability null is true, (R_t) converges to zero almost surely because it is dominated by the likelihood ratio corresponding to two singular distributions on the sequence space. \square

Example 9. Here we construct a non-exchangeable 2nd-order Markov chain (X_t) where the above test does not have power one, i.e., (R_t) does not converge to infinity, but instead to zero. This chain satisfies $p = 1/2$ and $\alpha = \gamma$, hence does not yield a counterexample to Theorem 7 but illustrates the importance of the condition $\alpha(1-p) \neq \gamma p$. To this end, set $p_{1|(0,0)} = p_{1|(0,1)} = 1$ and $p_{1|(1,0)} = p_{1|(1,1)} = 0$. To visualize this 2nd-order Markov chain, note that it alternates between two consecutive zeros and two consecutive ones. For large t we then have $n_{0|0} \approx n_{0|1} \approx n_{1|0} \approx n_{1|1} \approx t/4$ which yields for large t the bound

$$R_{4t} \leq \frac{\Gamma(t+1)^4}{\Gamma(2t)^2} = \frac{(t!)^4}{(2t-1)!^2}.$$

By Stirling's approximation of factorials, R_t tends to zero as t tends to infinity.

The above example motivates the following extension, which we only address briefly as a proof-of-concept since the details are conceptually straightforward but tedious.

2.5 Extension to higher-order Markovian alternatives and context trees

The aforementioned discussion derived a safe e-process and sequential test for exchangeability of a binary sequence, in particular targeting first-order Markovian alternatives. Neither the binary observations nor the first-order alternatives were particularly critical; these were adopted for clarity of exposition and simplicity of formulae.

Extensions to Markov sources of order $k > 1$ or alphabet sizes $d > 2$ are immediate. We may treat each k -th order context $x \in \{1, \dots, d\}^k$ as an independent d -ary prediction problem, and by mixing with independent Jeffreys' (which are Dirichlet(1/2, ..., 1/2)) priors (or equivalently, composing independent Krichevsky-Trofimov estimators), we obtain a computationally attractive e-process with regret bounded by $(d^k(d-1)/2) \ln t + O(1)$. In other words, we get a closed-form e-process $R_t^{k,d}$ — whose details are tedious, despite being explicit, and thus omitted — such that

$$R_t^{k,d} \geq \frac{\text{maximum likelihood of order } k \text{ Markov model}}{\text{maximum likelihood of Bernoulli model}} \cdot \exp \left(-\frac{d^k(d-1)}{2} \ln t - O(1) \right).$$

The (near)-optimality of this approach is discussed by Takeuchi et al. [33]. The e-process (R_t) from (6) can be interpreted as $(R_t^{1,2})$.

Further computationally attractive extensions include alternatives that consist of Markov sources of varying orders $k = 1, 2, \dots$ (see the discussion on the mixture method for unions below). The even more general Context Tree models have the length of the context that should be taken into account depend on that very context [45].

A similar calculation to the $k = 1, d = 2$ case done previously shows that $R_t^{k,d} \rightarrow \infty$, \mathbb{P} -almost surely for any alternative $\mathbb{P} \in \mathcal{P}_k \setminus \mathcal{Q}$, where \mathcal{P}_k is the set of Markovian distributions with order at most k . These developments lead naturally to the following subsection.

2.6 Double-mixtures for a countable sequence of alternatives

Let $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3 \dots$ be a countable sequence of alternatives, that may or may not be nested. Suppose for every $k \in \mathbb{N}$ one can design a safe e-process (E_t^k) for testing \mathcal{Q} against \mathcal{P}_k such that it has power one, meaning that for any $P \in \mathcal{P}_k \setminus \mathcal{Q}$, we have

$$\limsup_{t \rightarrow \infty} E_t^k = \infty, \text{ P-almost surely.}$$

Then, one can design a safe e-process for \mathcal{Q} against $\bigcup_{k \in \mathbb{N}} \mathcal{P}_k$ such that for any $P \in \bigcup_{k \in \mathbb{N}} \mathcal{P}_k \setminus \mathcal{Q}$, we have

$$\limsup_{t \rightarrow \infty} E_t = \infty, \text{ P-almost surely.}$$

The proof of the above claim is simple. We can, for example, define the “double mixture”

$$E_t := \sum_{k=1}^{\infty} \frac{6}{\pi^2 k^2} E_t^k,$$

which is a countable mixture over the base e-processes (that might have been already mixed using Jeffreys’ prior). It is straightforward to check that (E_t) is a safe e-process under \mathcal{Q} , by invoking monotone convergence and linearity of expectation. To analyze its power, once an alternative P has been picked, let \mathcal{P}_{k^*} be the first element of the nested sequence that contains P . Since $\limsup_{t \rightarrow \infty} E_t^{k^*} = \infty$, P-almost surely, the same property holds for $(E_t^{k^*}/k^{*2})$, and thus transfers to (E_t) since e-processes are nonnegative. The computational challenge of calculating (E_t) remains, but this can be reduced by instead calculating the \mathcal{Q} -safe e-process

$$\tilde{E}_t := \sum_{k=1}^t \frac{6}{\pi^2 k^2} E_t^k.$$

At the (finite) time k^* , \tilde{E}_t begins to include to required term $E_t^{k^*}$, and thus inherits its property of approaching infinity almost surely (consistency). Replacing the sum $\sum_{k=1}^t$ by $\sum_{k=1}^{f(t)}$ for any increasing function f that grows to infinity, possibly with sublinear growth (such as $\ln(\cdot)$), can further save computation without losing the consistency property.

2.7 Handling generic alternatives via non-anticipating likelihoods and betting

If the alternative to exchangeability is not clearly specified, then a different approach to the one above may be more suitable, as inspired by the recent work on universal inference by Wasserman et al. [42]. It involves a non-anticipating likelihood in the numerator, combined with an MLE in the denominator:

$$R_t^{\text{NA}} := \frac{\text{non-anticipating likelihood under the alternative}}{\text{maximum likelihood under the null}}.$$

Here, the numerator is simply given by

$$\prod_{s=1}^t g_s(X_s),$$

where g_s is any “non-anticipating” probability mass function, meaning that it is specified before seeing X_s , but can be learnt using the first $s - 1$ data points; in other words, (g_t) is predictable with respect to (\mathcal{F}_t) . One example would be to choose g_s as the (smoothed) maximum likelihood estimator under the alternative using the first $s - 1$ samples, but other approaches inspired by machine learning or time series modeling may also be employed. Note that any Bayesian mixture (including ours with Jeffreys’ prior) is of this non-anticipating form, with g_s being the associated predictive distribution.

It is easy to prove that (R_t^{NA}) is a \mathcal{Q} -safe e-process: each term can be verified to have conditional mean at most one by swapping the denominator for the (unknown) true null likelihood. The major strength of the above approach is that arbitrarily flexible nonparametric or model-free update rules can be used

without sacrificing validity, thus opening up the potential for power against loosely specified alternatives or even the discovery of temporal patterns from the observed data. For example, one may employ a complex Bayesian working model that outputs the posterior predictive probability of observing a zero or one at the next step, and this would not violate any of our theoretical guarantees regardless of the choice of priors or working model. Despite such a strong validity guarantee, the current drawback of this approach is that for generic update rules, there may not be an existing regret bound that we may use to convince ourselves of its power. (Of course, such regret bounds would be available for specific update rules and specific alternatives, and the online learning literature is rapidly expanding the scope and types of available regret bounds for individual sequence prediction.)

As a final remark, this non-anticipating likelihood is closely related to the “predictable-mixture” approach recently explored by [43], and has its roots in Wald [40, Eq. 10:10]. In this vein, it is also closely related to testing hypotheses by betting, as popularized by Shafer and Vovk [28, 30]; specifically (g_t) can be viewed as a sequence of bets on the following outcome.

3 Fork-convexity and \mathcal{Q} -Snell envelopes

Forgetting for a moment some of the earlier claims made without proof, one of the main questions we seek to answer in this section is:

When is a \mathcal{Q} -safe e -value simply a \mathcal{Q} -NM or \mathcal{Q} -NSM in disguise? In other words, is any \mathcal{Q} -safe e -value always improved (or recovered) by some \mathcal{Q} -NM or \mathcal{Q} -NSM?

Such a question was also asked in the latest preprint on safe testing by Grünwald et al. [10]. The necessity and sufficiency results of Ramdas et al. [22] imply that the answer in the singleton $\mathcal{Q} = \{\mathbf{Q}\}$ case is: *always* (via the Doob decomposition of the Snell envelope). The answer in the composite setting is: *sometimes*. We now qualify the ‘sometimes’ by delving into the rich probabilistic structure underlying safe e -values, examining its relationship to convex null sets, a concept called ‘fork-convexity’, and a process that we call a ‘composite’ Snell envelope, known from the mathematical theory of risk measures [6].

Most of this section does not depend on our observations being binary, and we allow the data (X_t) to take values in a more general space \mathcal{X} . Some of the technical notions required below, such as local absolute continuity, likelihood ratio (or density) processes, and essential suprema, are reviewed in Appendix A.

3.1 A sequential analog of convexity

We first review the concept of *fork-convexity*, which can be viewed as a sequential version of convexity.

Definition 10. Fix a reference measure \mathbf{R} on the sequence space $\mathcal{X}^{\mathbb{N}}$.

1. A fork-convex combination of two locally dominated laws \mathbf{Q}, \mathbf{Q}' with likelihood ratio processes $(Z_t), (Z'_t)$ is another law \mathbf{Q}'' with likelihood ratio process

$$Z_t'' := \begin{cases} Z_t, & t \leq s \\ hZ_t + (1-h)Z_s \frac{Z'_t}{Z'_s}, & t > s \end{cases} \quad (13)$$

for some $s \in \mathbb{N}_0$ and some \mathcal{F}_s -measurable random variable h in $[0, 1]$ with $h = 1$ on $\{Z'_s = 0\}$. The latter condition ensures that (Z_t'') is well-defined and an \mathbf{R} -martingale, as required for a likelihood ratio process.

2. A set \mathcal{Q} of probability measures is called fork-convex if every fork-convex combination of elements of \mathcal{Q} still belongs to \mathcal{Q} .

Fork-convexity was first introduced by Žitković [38]. It is closely related to a concept in the literature on risk measures called *m-stability*, due to Delbaen [6]. A similar notion called *rectangularity* was introduced by Epstein and Schneider [8] to describe intertemporal preferences with multiple priors. Rectangularity

has then been used extensively in the operations research literature in connection with robust Markov decision processes; see e.g. [14, 44, 32].

Note that fork-convexity implies convexity. To see this, observe that any (usual) convex combination $a\mathbb{Q} + (1-a)\mathbb{Q}'$ is also a fork-convex combination; just take $s = 0$ and $h = a$ in (13) to get $Z_t'' = aZ_t + (1-a)Z_t'$, which is the likelihood ratio process of $\mathbb{Q}'' := a\mathbb{Q} + (1-a)\mathbb{Q}'$.

A set $\{\mathbb{Q}\}$ that consists of a single law is clearly fork-convex. A set $\{\mathbb{Q}^1, \mathbb{Q}^2\}$ consisting of two distinct laws will not be fork-convex; it is not even convex. However, one can form its “fork-convex hull”. Here is the general definition.

Definition 11 (The fork-convex hull and its closure).

1. *The intersection of all fork-convex sets that contain a given set \mathcal{Q}_0 is called the fork-convex hull of \mathcal{Q}_0 . (Note that there is at least one fork-convex set containing \mathcal{Q}_0 , namely the set of all laws.)*
2. *The closed fork-convex hull of \mathcal{Q} is the closure of the fork-convex hull of \mathcal{Q} with respect to $L^1(\mathbb{R})$ convergence of the likelihood ratio processes at each fixed time $t \in \mathbb{N}$, where we recall \mathbb{R} is the assumed reference measure.*

Just as for usual convex hulls, the fork-convex hull of \mathcal{Q}_0 consists of all finite fork-convex combinations of elements in \mathcal{Q}_0 . Here a *finite fork-convex combination* of some distributions $\mathbb{Q}^1, \dots, \mathbb{Q}^n \in \mathcal{Q}_0$ is a distribution obtained by iteratively performing (13) a finite number of times on $\mathbb{Q}^1, \dots, \mathbb{Q}^n$, on their fork-convex combinations, on *their* fork-convex combinations, and so on. Closed fork-convex hulls play an important role in Theorem 13 below.

To provide some intuition for the definitions as applied to the null \mathcal{Q} considered in this paper, one can imagine a more “algorithmic” process of producing distributions in the closed fork-convex hull. First pick any $p_1 \in [0, 1]$ and observe $X_1 \sim \text{Ber}(p_1)$. Then, after observing X_1 , pick any p_2 , and observe $X_2 \sim \text{Ber}(p_2)$. Continue this process indefinitely. Then, the sequence (p_t) is (\mathcal{F}_t) -predictable and the resulting binary sequence has a law that is contained in the closed fork-convex hull of \mathcal{Q} .

It may be instructive to consider another simple example. For a fixed $\mu \in [0, 1]$, define \mathcal{Q}^μ as the set of product distributions \mathbb{Q} over infinite $[0, 1]$ -valued sequences (X_t) such that $\mathbb{E}_{\mathbb{Q}}[X_t | \mathcal{F}_{t-1}] = \mathbb{E}_{\mathbb{Q}}[X_t] = \mu$, and define $\tilde{\mathcal{Q}}^\mu$ as the set of distributions \mathbb{Q} (not necessarily of product form) over infinite $[0, 1]$ -valued sequences (X_t) such that $\mathbb{E}_{\mathbb{Q}}[X_t | \mathcal{F}_{t-1}] = \mu$. Then \mathcal{Q}^μ is not fork-convex if $\mu \in (0, 1)$ but $\tilde{\mathcal{Q}}^\mu$ is, and the latter is the closed fork-convex hull of the former. The problem of sequentially estimating μ in this setup has been recently studied by Waudby-Smith and Ramdas [43].

3.2 No power against fork-convex hulls

Consider a null set \mathcal{Q} locally dominated by a reference measure \mathbb{R} . We now establish the interesting fact that e-processes based on \mathcal{Q} -NSMs are powerless against any alternative in the closed fork-convex hull of \mathcal{Q} . We state this formally in Theorem 13 below, but the underlying reason is contained in the following lemma.

Lemma 12. *If (L_t) is a supermartingale under two laws \mathbb{Q}, \mathbb{Q}' , then (L_t) is also a supermartingale under every fork-convex combination \mathbb{Q}'' of \mathbb{Q} and \mathbb{Q}' .*

Proof. Note that \mathbb{Q}, \mathbb{Q}' are dominated by $\mathbb{R} := (\mathbb{Q} + \mathbb{Q}')/2$. Fix any $s \in \mathbb{N}_0$ and \mathcal{F}_s -measurable random variable h in $[0, 1]$, and let \mathbb{Q}'' be the fork-convex combination of \mathbb{Q}, \mathbb{Q}' given in (13). In compliance with the definition, we restrict h to satisfy $h = 1$ on $\{Z_s' = 0\}$. Suppose (L_t) is a supermartingale under \mathbb{Q} and \mathbb{Q}' . Equivalently, $(Z_t L_t)$ and $(Z_t' L_t)$ are supermartingales under \mathbb{R} . Thus for $t \in \{1, \dots, s\}$ we have

$$\mathbb{E}_{\mathbb{R}}[Z_t'' L_t | \mathcal{F}_{t-1}] = \mathbb{E}_{\mathbb{R}}[Z_t L_t | \mathcal{F}_{t-1}] \leq Z_{t-1} L_{t-1} = Z_{t-1}'' L_{t-1}.$$

For $t \geq s + 1$ we have

$$\mathbb{E}_{\mathbb{R}}[Z_t'' L_t | \mathcal{F}_{t-1}] = h \mathbb{E}_{\mathbb{R}}[Z_t L_t | \mathcal{F}_{t-1}] + (1-h) Z_s \mathbb{E}_{\mathbb{R}} \left[\frac{Z_t'}{Z_s'} L_t \middle| \mathcal{F}_{t-1} \right] \leq Z_{t-1}'' L_{t-1}.$$

Thus $(Z_t'' L_t)$ is an \mathbb{R} -supermartingale, or equivalently, (L_t) is a \mathbb{Q}'' -supermartingale. \square

The following theorem refers to the *closed* fork-convex hull of \mathcal{Q} . Recall that this is the closure of the fork-convex hull of \mathcal{Q} , understood in the sense of $L^1(\mathbb{R})$ convergence of the likelihood ratio processes at each fixed time $t \in \mathbb{N}$.

Theorem 13. *Let $\tilde{\mathcal{Q}}$ be the closed fork-convex hull of \mathcal{Q} . Then every \mathcal{Q} -NSM is in fact a $\tilde{\mathcal{Q}}$ -NSM. Thus a test based on a \mathcal{Q} -NSM is powerless against $\tilde{\mathcal{Q}} \setminus \mathcal{Q}$.*

Proof. The fork-convex hull of \mathcal{Q} consists of all finite fork-convex combinations of elements of \mathcal{Q} . Therefore, thanks to Lemma 12, every \mathcal{Q} -NSM remains an NSM under every law \mathbb{Q} in the fork-convex hull of \mathcal{Q} . To extend this to the closure, pick any element $\mathbb{Q} \in \tilde{\mathcal{Q}}$. Then there is a sequence (\mathbb{Q}^n) in the fork-convex hull of \mathcal{Q} such that $\mathbb{Q}^n \rightarrow \mathbb{Q}$. This means that $Z_t^n \rightarrow Z_t$ in $L^1(\mathbb{R})$ for all $t \in \mathbb{N}$, where (Z_t^n) and (Z_t) are the likelihood ratio processes of \mathbb{Q}^n and \mathbb{Q} , respectively. By passing to a subsequence, we may assume that $Z_t^n \rightarrow Z_t$, \mathbb{R} -almost surely, for all $t \in \mathbb{N}$. Let (L_t) be any \mathcal{Q} -NSM and hence a \mathbb{Q}^n -NSM for all n . Equivalently, $(Z_t^n L_t)$ is an \mathbb{R} -NSM for all n . By the \mathbb{R} -supermartingale property and the conditional version of Fatou's lemma, we get

$$\mathbb{E}_{\mathbb{R}}[Z_t L_t \mid \mathcal{F}_{t-1}] = \mathbb{E}_{\mathbb{R}} \left[\lim_{n \rightarrow \infty} Z_t^n L_t \mid \mathcal{F}_{t-1} \right] \leq \liminf_{n \rightarrow \infty} \mathbb{E}_{\mathbb{R}}[Z_t^n L_t \mid \mathcal{F}_{t-1}] \leq \liminf_{n \rightarrow \infty} Z_{t-1}^n L_{t-1} = Z_{t-1} L_{t-1}.$$

This completes the proof that every \mathcal{Q} -NSM is in fact a $\tilde{\mathcal{Q}}$ -NSM. \square

The first part of the above theorem asserts that the NSM property is preserved under taking closed fork-convex hulls, but note that this is not true for safe e -values in general. Indeed, (E_t) being \mathcal{Q} -safe implies that it is $\text{conv}(\mathcal{Q})$ -safe, but not necessarily $\tilde{\mathcal{Q}}$ -safe.

3.3 Composite Snell envelopes

For a single law $\mathbb{Q} \in \mathcal{Q}$ and an e -process (E_t) , the \mathbb{Q} -Snell envelope is the smallest \mathbb{Q} -NSM that dominates (E_t) . It is natural to ask whether, in contrast to this pointwise construction, one can directly construct a ‘‘composite \mathcal{Q} -Snell envelope’’, i.e., a smallest \mathcal{Q} -NSM that dominates (E_t) .

It turns out that the ability to define such a \mathcal{Q} -Snell envelope of an e -process depends heavily on the property of fork-convexity. The following result states that if the null set \mathcal{Q} is locally dominated and fork-convex, then a \mathcal{Q} -Snell envelope of a given e -process (E_t) exists, is safe, and improves upon (E_t) .

Theorem 14. *Let \mathcal{Q} be locally dominated and fork-convex. Let (E_t) be a \mathcal{Q} -safe e -process. Then the process*

$$L_t := \text{ess sup}_{\mathbb{Q} \in \mathcal{Q}, \tau \geq t} \mathbb{E}_{\mathbb{Q}}[E_{\tau} \mid \mathcal{F}_t], \quad t \in \mathbb{N}_0,$$

where τ ranges over all finite stopping times, is the smallest \mathcal{Q} -NSM that dominates (E_t) and satisfies $L_0 \leq 1$. Hence, (L_t) is the \mathcal{Q} -Snell envelope of (E_t) . In particular, by the optional stopping theorem, (L_t) is a \mathcal{Q} -safe e -process.

In this theorem, L_t is defined as the *essential supremum* of the family of random variables $\mathbb{E}_{\mathbb{Q}}[E_{\tau} \mid \mathcal{F}_t]$ indexed by τ and \mathbb{Q} . This means that L_t is the smallest random variable that almost surely dominates this family. The details of this definition are reviewed in Appendix A, along with some key properties. In particular, the proof of the theorem in Appendix B will make use of Proposition 20.

Remark 15. *We caution the reader that ‘essential supremum’ is used outside of this work with a different meaning: the essential supremum of a (single) random variable X is defined as the smallest constant c such that $X \leq c$ almost surely. This notion is different from the one used here, and does not arise in this paper.*

The process (L_t) above is what we call the \mathcal{Q} -Snell envelope. Note that the \mathcal{Q} -Snell envelope of (L_t) is almost surely equal to (L_t) itself. In short, the above theorem claims that if \mathcal{Q} is fork-convex, then the \mathcal{Q} -Snell envelope of any \mathcal{Q} -safe e -process exists and is safe.

To construct a powerful and valid test that dominates a safe e -process (E_t) , one might be inherently interested in the *largest* \mathcal{Q} -NSM (\bar{L}_t) that dominates (E_t) and satisfies $\bar{L}_0 \leq 1$. However, we are not aware

of a systematic way to obtain such a process. Nevertheless, even the smallest \mathcal{Q} -NSM that dominates (E_t) , namely the \mathcal{Q} -Snell envelope, still tends to improve its power.

For a given \mathcal{Q} , can there be more than one process that is considered a \mathcal{Q} -Snell envelope (of some other process), and amongst these, is there a largest one? In general, the answer is yes for the first question and (typically) no for the second. Every \mathcal{Q} -NSM is its own Snell envelope and there always exist uncountably many \mathcal{Q} -NSMs, namely the constant and nonnegative decreasing processes starting at one. In particular, the constant process is also a \mathcal{Q} -NM albeit a powerless one. In fact, there may be uncountably many \mathcal{Q} -NSMs, with none of these processes dominating the others, and at the same time there may not exist any non-constant \mathcal{Q} -NMs (that don't use independent external randomization, which involves expanding the filtration). For this paper's choice of \mathcal{Q} , we later show that every \mathcal{Q} -NSM is almost surely nonincreasing, and hence the constant process equaling one dominates all \mathcal{Q} -NSMs, and indeed the only \mathcal{Q} -NM almost surely equals one.

Taken together, Theorems 13 and 14 lead to the following corollary, which tells us that in certain situations one has to move beyond composite NSMs to achieve powerful tests. We continue to let \mathcal{Q} be any locally dominated null set and $\tilde{\mathcal{Q}}$ its closed fork-convex hull.

Corollary 16. *Let (E_t) be a \mathcal{Q} -safe e -process. Then (E_t) is dominated by (or equals) some \mathcal{Q} -NSM (L_t) with $L_0 \leq 1$ if and only if (E_t) already happens to be $\tilde{\mathcal{Q}}$ -safe (and therefore powerless against $\tilde{\mathcal{Q}} \setminus \mathcal{Q}$).*

Proof. To prove the forward implication, assume (E_t) is dominated by some \mathcal{Q} -NSM (L_t) with $L_0 \leq 1$. By Theorem 13, (L_t) is in fact a $\tilde{\mathcal{Q}}$ -NSM. It follows that (E_t) is $\tilde{\mathcal{Q}}$ -safe as claimed, because we have $\mathbb{E}_P[E_\tau] \leq \mathbb{E}_P[L_\tau] \leq L_0 \leq 1$ for every $P \in \tilde{\mathcal{Q}}$ and each finite stopping time τ .

To prove the reverse implication, assume that (E_t) is actually $\tilde{\mathcal{Q}}$ -safe. An application of Theorem 14 (with \mathcal{Q} replaced by $\tilde{\mathcal{Q}}$) then gives a $\tilde{\mathcal{Q}}$ -NSM (L_t) with $L_0 \leq 1$ that dominates (E_t) . This completes the proof of the corollary. \square

The above result suggests that we *must* look beyond NSMs for designing sequential tests for exchangeability, and we next show that this fact holds regardless of the class of alternatives considered.

3.4 The inadequacy of NSMs for testing exchangeability

We now return to the main focus of this paper, which is binary sequences; thus $\mathcal{X} = \{0, 1\}$. In this case, any law P is locally dominated by the i.i.d. Bernoulli(1/2) law $R := \text{Ber}(1/2)^\infty$ and the likelihood ratio process of P is

$$Z_t = 2^t \prod_{i=1}^t D_i, \quad \text{where} \quad D_i := q_i(X_1, \dots, X_{i-1}) \mathbf{1}_{\{X_i=1\}} + (1 - q_i(X_1, \dots, X_{i-1})) \mathbf{1}_{\{X_i=0\}} \quad (14)$$

for some functions $q_t: \{0, 1\}^{t-1} \rightarrow [0, 1]$ such that, R -almost surely,

$$q_t(X_1, \dots, X_{t-1}) = \mathbb{Q}(X_t = 1 \mid X_1, \dots, X_{t-1}).$$

In particular, taking $q_t = p \in (0, 1)$ for all t gives the likelihood ratio process of $\text{Ber}(p)^\infty$ with respect to R . The following theorem shows that any likelihood ratio process of the form (14) can be approximated by a finite fork-convex combination of likelihood ratio processes corresponding to i.i.d. Bernoulli laws. The closed fork-convex hull of this set is therefore very large: it contains *all* laws over binary sequences.

Theorem 17. *Every law P over the space of binary sequences belongs to the closed fork-convex hull of $\mathcal{Q} = \{\text{Ber}(p)^\infty : p \in (0, 1)\}$. Thus, the process (M_t) with $M_t = 1$ for all t is the only \mathcal{Q} -NM, and every \mathcal{Q} -NSM must be nonincreasing, and hence these never exceed one and always have zero power for any $\alpha \in (0, 1)$.*

Proof. Fix an arbitrary law P and let (Z_t) be its likelihood ratio process as in (14). We will show by induction that for each $s \in \mathbb{N}_0$ there exists an element $\mathbb{Q}^{(s)}$ in the fork-convex hull of \mathcal{Q} whose likelihood ratio process $(Z_t^{(s)})$ satisfies $Z_t^{(s)} = Z_t$ for all $t \leq s$. This is clearly true for $s = 0$ since every likelihood




Figure 2: A summary of some of the implications related to \mathcal{Q} -NSMs and \mathcal{Q} -safety (recall Figure 1 for the definitions of the classes \mathcal{Q} , $\bar{\mathcal{Q}}$, \mathcal{P}_1 , \mathcal{P}_k , and $\bar{\mathcal{Q}}$). We would like to design a \mathcal{Q} -safe e-process that is powerful against \mathcal{P}_k . Theorem 17 proves that a \mathcal{Q} -NSM is non-viable since it unintentionally results in \mathcal{P}_k -safety and thus no power against \mathcal{P}_k . The single non-implication sign above opens a door to constructing a non-NSM based \mathcal{Q} -safe e-process that is consistent against \mathcal{P}_k . An example for such an e-process is precisely the construction in (6).

ratio process is equal to one at time zero. Suppose it is true for some particular s . Consider a law \mathbb{Q}' whose likelihood ratio process (Z'_t) satisfies $Z'_{s+1}/Z'_s = 2D_{s+1}$, where D_{s+1} is defined in (14). We can choose \mathbb{Q}' from the fork-convex hull of \mathcal{Q} by mixing the 2^s Bernoulli laws $\text{Ber}(q_{s+1}(y_k))^\infty$ with weights $\mathbf{1}_{\{y_k\}}(X_1, \dots, X_s)$ at time s , where y_1, \dots, y_{2^s} lists all binary strings of length s . Taking the fork-convex combination of $\mathbb{Q}^{(s)}$ and \mathbb{Q}' at time s with $h = 0$ gives a law $\mathbb{Q}^{(s+1)}$ in the fork-convex hull of \mathcal{Q} whose likelihood ratio process satisfies $Z_t^{(s+1)} = Z_t^{(s)} = Z_t$ for $t \leq s$, and $Z_{s+1}^{(s+1)} = 2Z_s^{(s)}D_{s+1} = Z_{s+1}$ for $t = s + 1$. Thus $\mathbb{Q}^{(s+1)}$ satisfies the induction assumption with $s + 1$ instead of s .

Thus, by induction, for each $s \in \mathbb{N}_0$ there exists an element $\mathbb{Q}^{(s)}$ in the fork-convex hull of \mathcal{Q} whose likelihood ratio process $(Z_t^{(s)})$ satisfies $Z_t^{(s)} = Z_t$ for all $t \leq s$, as required. Now, it is clear that for each fixed $t \in \mathbb{N}$, $Z_t^{(s)} \rightarrow Z_t$ in $L^1(\mathbb{R})$ as $s \rightarrow \infty$. This shows that \mathbb{P} belongs to the closed fork-convex hull of \mathcal{Q} as claimed in the first part of the theorem.

It remains to prove that every \mathcal{Q} -NSM (L_t) must be nonincreasing, which implies in particular that the only \mathcal{Q} -NM is the constant process equal to one. Fix any binary sequence ω_0 and let \mathbb{P} be the law such that $\mathbb{P}(\{\omega_0\}) = 1$. By the first part of the theorem, \mathbb{P} belongs to the closed fork-convex hull of \mathcal{Q} and hence, by Theorem 13, (L_t) is a \mathbb{P} -NSM. Thus $L_t(\omega_0) = \mathbb{E}_{\mathbb{P}}[L_t] \leq \mathbb{E}_{\mathbb{P}}[L_{t-1}] = L_{t-1}(\omega_0)$ for every $t \in \mathbb{N}$. Since ω_0 was arbitrary, this shows that (L_t) must be nonincreasing. \square

In other words, not only are \mathcal{Q} -NSMs inadequate against Markovian alternatives, they are incapable of detecting *any* deviation from exchangeability. We include an overview of the relations we have shown in Figure 2.

4 Numerical simulations and some extensions

Recall that it is impossible to detect *all* deviations from exchangeability. For example, if only the first bit of the data is a $\text{Ber}(0.1)$ and the rest is $\text{Ber}(0.2)$, there is simply not enough information available to reject the null. The informal reason is that the deviation from the null is temporary and fleeting, not sustained,

and thus easily ascribed to chance. Motivated by this, the primary focus of the paper has been to detect Markovian alternatives. In this section, we will examine the power of our approach against three types of alternatives: (a) the one it was designed primarily for — a first- (and low-) order Markov alternative, (b) a time-inhomogeneous Markov alternative, and (c) a completely different type of alternative that we were not intending to have power against: a change point alternative.

Remark 18. *The last setting should especially not be confused with change point detection, which has different goals, such as minimizing ‘detection delay’ (informally, expected difference between the stopping time and the changepoint, when the procedure stops after the changepoint) subject to a prespecified lower bound on the ‘average run length’ (expected stopping time when there is no changepoint). To be clearer, in order for change point detection algorithms to have a small detection delay, they must pay the price of stopping under the ‘null’ as well, and one typically studies the tradeoffs between those two metrics. However, we will remain still in the realm of sequential testing, where we do not wish to ever stop if the null is true, and the metrics of Type I error and power are rather different from the above. Below, we will focus on measuring evidence using e-processes, where we wish to maximize wealth (or its rate of growth) under the alternative, subject to an upper bound on expected wealth under the null.*

We also reinterpret our test in terms of sequential estimation via confidence sequences in Section 4.6, and Section 4.7 demonstrates the use of calibrators to construct an e-process that approximately tracks the maximum wealth thus far.

4.1 Sanity check 1: no power against an i.i.d. Bernoulli sequence

We start with a sanity check, namely that our evidence measure does not report any evidence against Bernoulli sources (which are exchangeable). Theorem 7 predicts that when the null is true, (R_t) should decay to zero almost surely, meaning that $(\ln R_t)$ decays to minus infinity. Further, Jensen’s inequality implies that at any stopping time τ , $\ln R_\tau$ has nonpositive expectation; hence $\sup_{\mathcal{Q} \in \mathcal{Q}} \sup_{\tau} \mathbb{E}_{\mathcal{Q}}[\ln R_\tau] \leq 0$, where the second supremum is over all stopping times τ . Figure 3 illustrates that indeed the evidence decays from the start, as expected. The lower bound from (8) predicts that $\ln R_t$ could decay logarithmically in time (but not faster), meaning that $\ln R_t \gtrsim -\ln t$, and that seems particularly tight in the figure below.




Figure 3: The blue line shows the evolution of $\ln R_t$ under two Bernoulli sources. The green line is the lower bound $-\ln(t)$ following from (8). The decay is consistent with Theorem 7, which predicts that $\ln R_t \rightarrow -\infty$ almost surely.

Having established that the predictions under the null are accurate, we now estimate the power against Markovian alternatives. We first consider a fixed Markov alternative, which the test targets. Then we consider a time-varying Markov alternative, which the test does not explicitly target.

4.2 Sanity check 2: power against a first-order Markov alternative

Here we evaluate the evidence measure (6) in the well-specified case where data come from a first-order Markov process. Figure 4 illustrates that the asymptotics given by Theorem 7 kick in early, as evidenced by the (log)-evidence growing linearly with the sample size.




Figure 4: Blue: the evolution of $\ln R_t$, under a Markov process that is relatively far from Bernoulli (left) and one that is closer (right). Red: the linear approximation with slope r^* from Theorem 7.

4.3 Power against a time-varying Markov alternative

Recalling (9), one might conjecture that if a sequence has equal zeroth and first-order frequency limits, then our test will be powerless. Here, we show that this is not necessarily the case. Namely, we show that it is possible for a sequence to have frequency limits $\lim_{t \rightarrow \infty} n_1/t = \lim_{t \rightarrow \infty} n_{1|0}/n_0 = \lim_{t \rightarrow \infty} n_{1|1}/n_1$ (this happens almost surely under a Bernoulli source) while (R_t) still diverges to infinity, meaning that we reject the exchangeability null at any confidence level. We will consider a time-varying Markov source with symmetric transition probabilities $p_{0|0}(t) = p_{1|1}(t) = 1/2 + \delta_t$, for some decreasing sequence δ_t tending to 0. To make the example interesting, we need two properties of the “repetition bias” δ_t :

- We aim to pick δ_t small enough so that the above limit frequencies are $1/2$. In particular, this requires $\lim_{t \rightarrow \infty} \sum_{s=1}^t \delta_s/t = 0$. For example, $\delta_t := 1/t^\alpha$ for $\alpha > 0$ would do for this purpose.
- We aim to pick δ_t large enough so that $\ln R_t \rightarrow \infty$. To this end, we can use the Law of the Iterated Logarithm to conclude that after t rounds n_0 and n_1 are both $t/2 \pm O(\sqrt{t \ln \ln t})$ while $n_{0|0}$ and $n_{1|1}$ are both $t/4 + \sum_{s=1}^t \delta_s/2 \pm O(\sqrt{t \ln \ln t})$. Let us write $\ell(t)$ for the likelihood ratio of the Markov maximum likelihood and the Bernoulli maximum likelihood, as we did in the proof of Theorem 7. Abbreviating the binary entropy to $h(p) := -p \ln p - (1-p) \ln(1-p)$ and using that $h(1/2 + \delta) = \ln 2 - 2\delta^2 + O(\delta^3)$ reveals that


$$\ell(t) = t \left(h\left(\frac{n_1}{t}\right) - \frac{n_0}{t} h\left(\frac{n_{0|0}}{n_0}\right) - \frac{n_1}{t} h\left(\frac{n_{1|1}}{n_1}\right) \right) \approx \frac{2}{t} \left(\sum_{s=1}^t \delta_s \pm O(\sqrt{t \ln \ln t}) \right)^2. \quad (15)$$

To have the lower bound $\ln R_t \geq \ell(t) - \ln t - O(1)$ from (8) diverge, we need to ensure

$$2 \left(\frac{1}{\sqrt{t}} \sum_{s=1}^t \delta_s - O(\sqrt{\ln \ln t}) \right)^2 - \ln t \rightarrow \infty.$$

This reveals that the threshold for divergence is around $\sum_{s=1}^t \delta_s = \sqrt{t/(2 \ln t)}$ (this being forced by the $\ln t$ term on the right, not by the $\sqrt{\ln \ln t}$ term inside the square). To investigate the

behaviour ever so slightly above this threshold, we ended up picking $\delta_t := F(t+1) - F(t)$ for $F(t) := \min\{t/2, \sqrt{t} \ln(1+t)\}$, where the min ensures that $\delta_t \leq 1/2$ from $t = 1$ without affecting the order of growth. The results are displayed and discussed in Figure 5.



(a) Markov($1/2 - \delta_t, 1/2 + \delta_t$). This source is “sticky”, in that it slightly favours repeating the previous outcome. The evidence against the null diverges. Red is the approximation (15) with the constant in $O(\cdot)$ taken to be 1, while green is (15) with the constant in $O(\cdot)$ as -1 and $\ln(t)$ subtracted.

(b) Bernoulli($1/2 + \delta_t$) = Markov($1/2 + \delta_t, 1/2 + \delta_t$). This source is one-sided, in that it slightly favours ones. No evidence against the null accumulates. The green line is the lower bound $-\ln(t)$ following from (8).

Figure 5: The evolution of our evidence ($\ln R_t$) for a time-varying Markov process. On the left the transition probabilities are $p_{1|1}(t) = p_{0|0}(t) = 1/2 + \delta_t$ for $\delta_t := F(t+1) - F(t)$ with $F(t) := \min\{t/2, \sqrt{t} \ln(1+t)\}$. On the right, we perform a sanity check with $p_{1|1}(t) = p_{1|0}(t) = 1/2 + \delta_t$ instead, which renders this a time-varying Bernoulli process. As predicted, on the left the log-evidence diverges, growing at a logarithmic rate (which is far slower than the linear rate of Figure 4). The right picture confirms that this is due to the Markovian aspect of the source, and not due to the changing probabilities.

Despite the generality and pathwise nature of Theorem 7, this experiment points towards the possibility of strengthening it even further. Indeed, Theorem 7 has a very natural sufficient condition for consistency, but this experiment shows that even when this condition does not hold (meaning that the limits (9) exist and the condition (10) fails), then our safe e-process can still be consistent in certain subtle settings. These subtleties were apparent to us when designing this example, but we are not yet sure how to formalize these observations in a general fashion, and thus leave this to future work.

Another such example that demonstrates the unexpected power of our approach is examined next.

4.4 Power against a change point alternative

We consider a somewhat counterintuitive example to show that a first-order Markov alternative to exchangeability is perhaps more powerful than one may believe at first sight. Consider a length $2n$ sequence of coin flips sampled from $\text{Ber}(p)^n \text{Ber}(q)^n$ for some $p \neq q$. To match the setup of this example with our initial problem set up, one could potentially extend this to an infinite sequence in an arbitrary way, for example just continuing as $\text{Ber}(q)$ after time $2n$.

This sequence is clearly not exchangeable. It is, however, not clear whether our proposed first-order Markov alternative would detect (much) evidence against the null, as the sequence is not Markov, but is more like a change point alternative. Detecting evidence is not a given; the outcomes are in fact independent (albeit not identically distributed). Hence there is no first-order dependency structure for the Markov model to exploit. And on top of that, there seems to be only one problematic time-point, precisely half-way through the sequence. So even if the Markov model somehow exploited this, how could it gain an amount of evidence growing with the length n of the sequence?

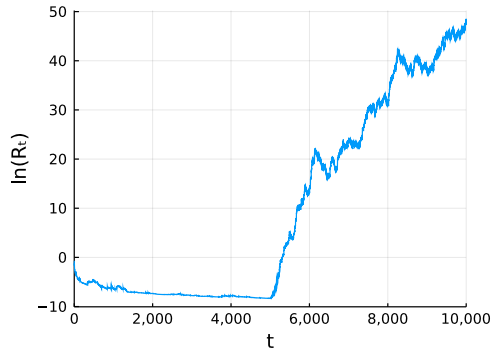


Figure 6: The process $(\ln R_t)$ on a sequence sampled from $\text{Ber}(.1)^{5000}\text{Ber}(.4)^{5000}$. On the first half, we see (R_t) decays as $1/t$, which is due to the overhead of Jeffreys’ mixture for the Markov model over the maximum likelihood Bernoulli parameter. After the change point, we see (R_t) increasing fast on the exponential scale. Recalling (4) for those more familiar with the p-value scale, the corresponding anytime p-process dips below 10^{-20} towards the end.

We now show that the above arguments are all misguided, and that the process (R_t) from (6) gains an amount of evidence against the exchangeable null that grows exponentially with t , between time n and $2n$. The evolution of (R_t) on a typical run of this process is shown in Figure 6.

Initially, (R_t) loses steam and tends towards zero at a rate $1/t$ before time n since the null is true and there is a price to pay for the Jeffreys’ mixture over the alternative. To calibrate what to expect after the change point, think of n as being relatively large so that we can reason about empirical frequencies of zeros and ones with more ease. Let us compute the maximum likelihood parameters for typical sequences with frequency (tending to) p in the first half and q in the second half. For the Bernoulli model, we find $\hat{p} = (p + q)/2$. For the first-order Markov model we find that

$$\hat{p}_{1|1} = \frac{p^2 + q^2}{p + q} \quad \text{and} \quad \hat{p}_{1|0} = \frac{(1 - p)p + (1 - q)q}{(1 - p) + (1 - q)}.$$

The main observation here is that the best Markov model is i.i.d. if $\hat{p}_{1|1} = \hat{p}_{1|0} = \hat{p}$, which occurs if and only if $p = q$. The fact that an exploitable first-order Markov dependency structure arises can perhaps be best observed in the extreme case $p = 0$ and $q = 1$. As this comparison does not really depend on n , we find that for all other parameter settings with $p \neq q$, the Markov model will gain overall evidence exponentially growing with t between time n and $2n$. (Technically, the exponential growth does not start immediately at time $n + 1$, but it does so eventually.) However, as t grows even further — say beyond $t = n^2$ or $t = 2^n$ — (R_t) will decrease once more towards zero. This is because the sequence eventually is dominated by i.i.d. $\text{Ber}(q)$ coin flips, and the MLE under the null explains the data very well.

Thus, for this example, we do not get a power one test, nor should we expect a single change point away from an i.i.d. model to yield power one for a test designed to be powerful against Markovian alternatives. However, if the initial pattern repeats itself after $2n$ steps, meaning that we keep alternating between $\text{Ber}(p)$ and $\text{Ber}(q)$ models, then (R_t) does have power one, and this is interesting because (R_t) is designed for first-order Markov alternatives, but it is consistent against these $2n$ -th order Markov alternatives.

In fact, one can argue that it is information theoretically impossible to design *any* power one test, including tests that are tuned to detect a single change point. To see why, think of very small n , like $n = 2$, to make the reason intuitively transparent. How can a test possibly have enough evidence with just one or two coin flips before the change point, to know with probability one a change actually did occur? Naturally, the larger the time n of the change point, the higher the power could be of any such test (as it is for our test also), but no test can possibly have power one since there is always some small probability (vanishing with n) that the distribution of the first n coin flips look quite similar to the post-change distribution.

Nevertheless, this simple example illustrates the point that our proposed e-process (R_t) for evidence against exchangeability is actually powerful even in scenarios that are not (close to) Markov.

Deriving other e-processes targeted towards change point alternatives. For readers explicitly interested in powerful tests to detect change point alternatives in the setting of this paper, we briefly describe a powerful test (albeit not a power one test, as already explained above). Essentially, one can combine the ideas in Remark 2.6 with those in Section 2.7. We let \mathcal{P}_k denote the alternative in which the change point is hypothesized to occur at time $n = 2^k$, though other increasing functions of k may also suffice. We will define an e-process (E_t^k) for each k and then use a countable mixture over k as the final e-process.

Now, we describe an e-process for a fixed k . Define $(g_s^-)_{s=1}^n$ to be a smoothed non-anticipating maximum likelihood estimator, calculated using data from time 1 to $s - 1$. The smoothing step is simple: add a single fake observation worth half a heads (or half a tails) to the counts when determining the MLE. The smoothing leads to a slight regularization that can be viewed as the maximum-a-posteriori estimate using a Beta(1/2, 1/2) prior, analogous to Krichevsky-Trofimov betting [20]. Similarly, define $(g_s^+)_{s=n+1}^\infty$ to be the same smoothed non-anticipating maximum likelihood estimator, but calculated using data from time $n + 1$ to $s - 1$. In both cases, the smoothing also leads to a well-defined function g_1^- and g_{n+1}^+ , which are effectively treated as a Ber(1/2) model. Finally, define the e-process (E_t^k) as the ratio of $\prod_{s=1}^{n \wedge t} g_s^-(X_s) \prod_{s=(n+1)}^t g_s^+(X_s)$ to the maximum likelihood under the i.i.d. null. In other words, the denominator is identical to one of R_t , but the numerator has changed because the targeted alternative is now different.

Recalling Section 2.7 (and the final section of Wasserman et al. [42]), it is easy to see that (E_t^k) is a \mathcal{Q} -safe e-process. If a change point occurs at time n^* (and let $k^* := \lfloor \ln_2 n^* \rfloor$), the e-processes $E_t^{k^*}$ and $E_t^{k^*+1}$ will grow exponentially between time n^* and $2n^*$. Even with the countable weighting of Remark 2.6, their exponential growth washes out the inverse polynomial weights, to yield a powerful e-process (E_t) .

Naturally, many permutations and combinations of these ideas (see e.g. the literature on switching experts [18]) can be used to derive a variety of tests against different kinds of alternatives. See Section 5.3 and [37] for an approach based on conformal prediction. We leave further exploration of these variants to future work.

4.5 Power on a real-world data set

We conclude the empirical evaluation with a jovial yet practical question: if we track the days on which it rains in London, do we get an exchangeable sequence? Three of us having lived there deemed this hypothesis unlikely yet not implausible, but a Markovian alternative certainly seems a-priori more plausible. To answer our question, we obtained daily areal rainfall data for North London from <https://data.london.gov.uk/dataset/daily-areal-rainfall>, spanning the 4404 days between 01/01/2007 to 21/01/2019. We binarised the data by comparing the daily total millimeters to zero, resulting in 59% rainy days overall. As is clear from Figure 7, the null hypothesis can be safely rejected.

4.6 An equivalent viewpoint based on confidence sequences

It is possible to view our sequential test — rejecting the null if (R_t) ever exceeds $1/\alpha$ — in terms of sequential estimation using confidence sequences [2, 13, 22]. Denoting $\mathbf{Q}_p := \text{Ber}(p)^\infty$, a $(1 - \alpha)$ -confidence sequence for p is a sequence of confidence sets $(C_t)_{t \geq 1}$ such that $\mathbf{Q}_p(\exists t \geq 1 : p \notin C_t) \leq \alpha$. Recalling (7), define

$$C_t := \left\{ q : R_t^{\text{JP},q} < \frac{1}{\alpha} \right\}.$$

To see that (C_t) is a confidence sequence for p , note that $p \notin C_t$ if and only if $R_t^{\text{JP},p} \geq 1/\alpha$, but Ville's inequality yields that $\mathbf{Q}_p(\exists t \geq 1 : R_t^{\text{JP},p} \geq 1/\alpha)$ since $R_t^{\text{JP},p}$ is a \mathbf{Q}_p -NM.

Based on the above argument, Theorems 6 and 7 together imply the following:




Figure 7: The process $(\ln R_t)$ accumulating evidence against the hypothesis that daily rain in North London is an exchangeable process.

- Rejecting the null as soon as (C_t) becomes empty is identical to rejecting if (R_t) ever exceeds $1/\alpha$:

$$\inf\{t \geq 1 : C_t = \emptyset\} = \inf\left\{t \geq 1 : R_t \geq \frac{1}{\alpha}\right\}. \quad (16)$$


This is because $R_t = \inf_{q \in [0,1]} R_t^{\text{JP},q}$, making the decision rules identical.

- If the data X_1, X_2, \dots are exchangeable, then with probability at least $1 - \alpha$, C_t is always nonempty. Indeed, it is possible to argue that C_t will shrink to a single point with probability $\geq 1 - \alpha$. Indeed, a single draw from $\mathbf{Q} \in \overline{\mathcal{Q}}$ is simply an i.i.d. $\text{Ber}(p)^\infty$ sequence for some p ; and for any $q \neq p$, $R_s^{\text{JP},q}$ will eventually exceed $1/\alpha$ because Jeffrey’s mixture likelihood (a mixture over all first-order Markov chains, including the special case that recovers an i.i.d. $\text{Ber}(p)$ sequence) will explain the data much better than the $\text{Ber}(q)$ likelihood. This can be formalized via the regret bound (8) but this simple argument is omitted for brevity.
- Since C_t contains the aforementioned p under the null (simultaneously for all t with probability at least $1 - \alpha$), we can infer that $\bigcap_{t=1}^\infty C_t \equiv \left\{q : \sup_{t \geq 1} R_t^{\text{JP},q} < 1/\alpha\right\}$ is also nonempty with probability at least $1 - \alpha$. This means that the decision rule $\inf\{t \geq 1 : \bigcap_{s \leq t} C_s = \emptyset\}$ yields a valid level α sequential test that could potentially halt sooner than (16). (The advantage is only expected to be slight in practice.)
- For any sequence of observations where the limits in (9) exist but (10) fails to hold, C_t will become empty at some point (deterministically). As explained around Theorem 7, the aforementioned failure of (10) happens with probability one under first-order Markovian alternatives, for example.

Figure 8 illustrates these points using the running intersection $\bigcap_{s \leq t} C_s$ on synthetic and real-world data.

4.7 Calibrated p-processes and adjusted e-processes for not losing capital

While (R_t) is a \mathcal{Q} -safe e-process, $(\max_{s \leq t} R_s)$ is not. In other words, we are only allowed to measure our performance based on the wealth accumulated thus far and not the highest wealth that we reached at some point in the process. The same is not true for p-processes: $(1/R_t)$ is a \mathcal{Q} -valid p-process, and so is $(1/\max_{s \leq t} R_s)$, the latter being the running infimum of the former. In game-theoretic terminology, the gambler can decide to stop playing the game (betting against the null) according to any stopping rule τ , but once they have stopped, only the final wealth R_τ of the gambler matters, and a nearly bankrupt gambler cannot point to their past wealth as a measure of their proficiency. This subtle point particularly manifests itself in the above change point example, because with a single change point, (R_t) rises to some amount (above 10^{20} in the figure) and then will shrink back to zero, so if we happen to stop too late, then R_τ could provide only meagre evidence even though it was once astronomically large.



(a) The rain dataset from Section 4.5. At all α considered, the confidence interval becomes empty, meaning that we reject the null. (b) $\text{Ber}(0.2)$ (as in Section 4.1). Here the null is true, and indeed all confidence intervals stay non-empty.

Figure 8: The confidence sequence (anytime-valid confidence interval) $\bigcap_{s \leq t} C_s$ for p is depicted by its lower and upper boundaries as a function of t , at various coverage probabilities $\alpha \in (0, 1)$. The running average of the data is displayed in black. As predicted, the confidence sequence becomes empty when the null is false (left), while it stays non-empty when the null is true (right).

So how can we get around this worrisome issue? We take inspiration from Shafer et al. [31] and use “calibrated p-processes” as our e-processes. (As a matter of terminology, our use of calibration here can be seen as an $E \rightarrow P \rightarrow E$ process, but if we skip the middle step entirely, the $E \rightarrow E$ direct method has been called “adjustment” by Dawid et al. [4, 3], Koolen and Vovk [19]. We will present it from both angles below to tie some loose ends in the literature together.)

Define $\mathbf{p}_t := 1/\max_{s \leq t} R_s$, so that (\mathbf{p}_t) is a \mathcal{Q} -valid p-process that satisfies (4). Let f be a calibrator [31, 22], which is a nonincreasing function f such that $\int_0^1 f(u)du = 1$. Then $(f(\mathbf{p}_t))$ is a \mathcal{Q} -safe e-process. It is not hard to check that $f(\mathbf{p}_t) \leq R_t$, so there is some price to pay for being able to take the best possible wealth into account. One possible choice for f is given by

$$f(u) := \frac{1 - u + u \ln u}{u(-\ln u)^2};$$

also see Vovk and Wang [36, Eq. (2)].

In order to do things more directly, let F be an adjuster [31, 3], which is an increasing function F such that $\int_1^\infty F(y)y^{-2}dy = 1$. Defining $A_t := F(\max_{s \leq t} R_s)$ yields that (A_t) is a \mathcal{Q} -safe e-process, and indeed as before $A_t \leq R_t$. One possible choice for F is given by

$$F(y) := \frac{y^2 \ln 2}{(1 + y)(\ln(1 + y))^2}.$$

Thus, even if (R_t) rises sharply and then decreases to zero eventually, (A_t) does not since it is nondecreasing. In fact, using the F given above in our example with a single change point, and noting that $F(y) \asymp y/(\ln y)^2$ for large y , we see that $A_\infty \approx 10^{17}$ even though $R_\infty = 0$. Of course, if $R_t \rightarrow \infty$ then so does $F(R_t)$, meaning that it does not lose the consistency property against Markovian alternatives.

Thus, at a (squared) logarithmic price to the overall capital, one can be protected against future losses, and for this reason we recommend using $A_t = F(R_t)$ as an e-process if we are uncertain about how close our alternative might be to the idealized Markovian case studied here.

5 Summary and discussion

5.1 From convex hulls to fork-convex hulls

The celebrated theorem of de Finetti — for which many proofs exist including based on elementary arguments [17] — states that all exchangeable binary sequences are mixtures of i.i.d. sequences (this was generalized much beyond the binary setting by Hewitt and Savage [11]). In fact, for any exchangeable sequence, the empirical measure $P_t := (1/t) \sum_{s=1}^t \delta_{X_s}$ converges in distribution to a measure μ supported on $[0, 1]$, and this is the so-called “de Finetti mixing measure” alluded to in the previous sentence. The crux of the matter is that the convex hull of all i.i.d. distributions over infinite binary sequences is precisely the set of exchangeable distributions. Since the convex hull preserves properties like safety (of e-processes) and validity (of sequential tests and p-processes), one can develop tests for the i.i.d. setting, which is itself a nontrivial composite null, and invoke de Finetti’s theorem to extend the result to the exchangeable setting.

In this paper, we go several steps further: we prove that the set of Markovian sequences lies inside the “fork-convex hull” of all exchangeable (or i.i.d.) sequences. In fact, Theorem 17 shows that the closed fork-convex hull is so large that *every* law over binary sequences is contained in it!

The fork-convex hull of a set of distributions can be informally thought of as a “predictable mixture of these distributions” (here, we borrow the terminology of predictable mixtures from Waudby-Smith and Ramdas [43], where it was applied to processes). Speaking informally in language inspired from game-theoretic probability, if Reality wanted to draw an infinite sequence from the fork-convex hull of \mathcal{Q} , it would pick \mathbb{Q}^1 and draw (only) X_1 from it, then based on that the outcome, Reality would pick \mathbb{Q}^2 and draw (only) X_2 from its conditional distribution given X_1 , and so forth ad infinitum, each time using the observed data to pick the next (conditional) distribution as it wishes.

Fork-convexity was a central object in this paper, and one of its primary roles was in dealing a fatal blow to tests based on constructing nonnegative supermartingales, as summarized next.

5.2 The powerlessness of test supermartingales, and a powerful e-process

Theorem 13 shows that the nonnegative supermartingale (NSM) property is preserved not just by taking the convex hull of a set of distributions, but also when taking the (much larger) fork-convex hull. Using composite Snell envelopes, Corollary 16 shows that if a safe test for \mathcal{Q} is upper bounded by a \mathcal{Q} -NSM, then it must have already been safe for the fork-convex hull of \mathcal{Q} (and thus powerless). Together, these results show that any NSM under exchangeable distributions is also an NSM under Markovian distributions, and in fact it is an NSM under *every* distribution over binary sequences. In other words, test statistics that are NSMs (or “test supermartingales”) are nonincreasing sequences, rendering them powerless to distinguish non-exchangeable distributions from exchangeable ones.

We get around the above hurdles by designing a safe e-process (R_t) in (6) that is upper bounded by some nonnegative martingale for every exchangeable distribution, despite not being an NSM itself. (The idea of designing a process that is upper bounded by an NSM, despite not being one itself, also appears elsewhere in the sequential testing literature [12], but for different reasons.) Our safe e-process uses the method of mixtures with Jeffreys’ prior to handle the composite alternative, along with the maximum likelihood under the null, to ultimately yield a computationally efficient closed-form e-process. This e-process not only has the desired safety properties at arbitrary stopping times (Theorem 6), but we prove that it has power one against any first-order Markovian alternative, and also a generic dense set of higher-order Markovian alternatives (Theorem 7). For first-order alternatives, the e-process also grows at an optimal rate, as implied by a regret bound (8) borrowed from the universal coding literature.

Section 2 also describes how to derive other e-processes that work for higher-order Markovian alternatives, and finally also for even more general, loosely specified alternatives by combining the method of predictable mixtures [43] and betting [28], along with universal inference [42].

5.3 Vovk’s approach based on conformal prediction

An alternate approach towards testing exchangeability was recently expounded by Vovk [35], which is based on conformal prediction. It replaces the canonical filtration (\mathcal{F}_t) by a poorer filtration (\mathcal{G}_t) formed

by conformal p-values. (We refer to their paper for the technical details.) Vovk then produces a sequence of independent p-variables under the null, which are converted to e-processes by appropriate calibration, which are in turn combined to form a martingale with respect to (\mathcal{G}_t) . This is particularly interesting because, despite the only martingales with respect to (\mathcal{F}_t) being constants, Vovk is able to identify nontrivial martingales with respect to an appropriately impoverished filtration (\mathcal{G}_t) .

Our approaches based on Jeffreys' mixture and the nonanticipating likelihood (or predictable mixture) can be seen as providing two alternatives to Vovk's methodology. Vovk's algorithm seems particularly powerful for change point alternatives, making them most similar to the extensions discussed in Section 4.4, while our paper focuses more on Markov alternatives. Further, since Vovk's martingale property only holds with respect to (\mathcal{G}_t) , so are the set of stopping times for which his method has the safety property. In other words, the algorithm is not allowed to look at the raw data and decide when to stop; it must only look at the sequence of p-variables. In contrast, our method is safe with respect to a much larger class of stopping times, indeed *all possible* stopping times (with respect to (\mathcal{F}_t) , the canonical filtration). However, in the end, the details of both works appear to be very different, and the conceptual principles by which the methods are derived also differ significantly.

A final, alternate approach to this problem could utilize *reverse* martingales and *exchangeable* filtrations. To elaborate, the exchangeable filtration is the reverse filtration $(\mathcal{E}_t)_{t=0}^\infty$ where $\mathcal{E}_0 := \sigma(\{X_1, X_2, \dots\})$, and for all $t \geq 1$, \mathcal{E}_t denotes the σ -algebra generated by all real-valued Borel-measurable functions $f(X_1, X_2, \dots)$ which are permutation-symmetric in their first t arguments, so that $\mathcal{E}_0 \supseteq \mathcal{E}_1 \supseteq \mathcal{E}_2 \dots$. It is known that if the data are exchangeable, then the empirical measure $P_t := (1/t) \sum_{s=1}^t \delta_{X_s}$ forms a measure-valued reverse martingale with respect to the exchangeable filtration, in the sense that $(\int g dP_t)$, is a reverse martingale for any bounded and Borel-measurable function g [15]. In fact, the converse of this statement also holds true if the sequence (X_t) is stationary [1]. This fact has recently been exploited to develop confidence sequences and sequential tests in other contexts [21]. We hope to explore in more detail whether this approach can lead to powerful tests for exchangeability in the future.

5.4 A game-theoretic protocol and reinterpretation of Theorem 7

Despite our work being grounded in measure theory, it has some implications for game theoretic probability, as developed by Shafer and Vovk, amongst others. Due to lack of space, we presume the reader has some general familiarity with their setup and refer to their latest book [30] for details. We recall from Shafer [28] the basic protocol of testing by betting:

- A skeptic begins with initial wealth of one, $\mathcal{K}_0 := 1$.
- for $t = 1, 2, \dots$
- A forecaster announces a distribution Q_t for X_t .
- A skeptic announces a bet $S_t : \{0, 1\} \rightarrow \mathbb{R}^+$ such that $\mathbb{E}_{Q_t}[S_t(X_t)] \leq 1$.
- Reality announces the outcome x_t .
- The skeptic's wealth gets updated as $\mathcal{K}_t := \mathcal{K}_{t-1} \cdot S_t(x_t)$.

At any point in time, the skeptic's wealth \mathcal{K}_t acts as a measure of evidence against the forecaster. Instead of announcing a single distribution Q_t , the forecaster could announce a set of distributions \mathcal{Q}_t , in which case the bet S_t must satisfy the constraint that $\sup_{Q_t \in \mathcal{Q}_t} \mathbb{E}_{Q_t}[S_t(X_t)] \leq 1$. While this protocol is very general, the setup for the current paper is better understood using the following protocol.

- A forecaster F announces a set of distributions $\overline{\mathcal{Q}}$ for the entire sequence.
- A skeptic S wishes to challenge F, and observes that $\overline{\mathcal{Q}} = \text{conv}(\bigcup_{m \in \mathcal{M}} \mathcal{Q}^m)$, for some set \mathcal{M} , where \mathcal{Q}^m is itself a set of distributions.
- S requests F to play a different game for each $m \in \mathcal{M}$, starting each with one dollar, $\mathcal{K}_0^m := 1$.
- F accepts the request, as long as the Skeptic's net wealth is measured by its worst performance across all games, that is, $\mathcal{K}_t := \inf_{m \in \mathcal{M}} \mathcal{K}_t^m$.
- for $t = 1, 2, \dots$
- For each $m \in \mathcal{M}$, Skeptic places a bet $S_t^m : \{0, 1\} \rightarrow \mathbb{R}^+$ such that $\sup_{Q_t^m \in \mathcal{Q}_t^m} \mathbb{E}_{Q_t^m}[S_t^m(X_t)] \leq 1$, where $Q_t^m(X_t) := Q^m(X_t | X^{t-1})$ for each $Q^m \in \mathcal{Q}^m$ and $\mathcal{Q}_t^m := \{Q_t^m : Q^m \in \mathcal{Q}^m\}$.

- Reality announces the (common across games) outcome x_t .
- The skeptic’s wealth in the m -th game gets updated as $\mathcal{K}_t^m := \mathcal{K}_{t-1}^m \cdot S_t^m(x_t)$.

In this paper, of course, $\mathcal{M} := [0, 1]$, $m := p$, and $\mathcal{Q}^m := \{\text{Ber}(p)^\infty\}$ is a singleton. The skeptic’s wealth $\mathcal{K}_t := \inf_{m \in \mathcal{M}} \mathcal{K}_t^m$ would be the game-theoretic instantiation of an e-process (while in the earlier setup, it would have been a supermartingale). Even though we do not employ the added generality, it is worth noting that both the above protocols are just special cases of the following protocol that affords the forecaster more flexibility.

- S and F agree to play a set of games indexed by $m \in \mathcal{M}$, starting each with one dollar, $\mathcal{K}_0^m := 1$. The Skeptic’s net wealth is measured by $\mathcal{K}_t := \inf_{m \in \mathcal{M}} \mathcal{K}_t^m$ for $t = 1, 2, \dots$
- For each $m \in \mathcal{M}$, Forecaster announces a set of distributions \mathcal{Q}_t^m on \mathcal{X}_t .
 - For each $m \in \mathcal{M}$, Skeptic places a bet $S_t^m : \mathcal{X}_t \rightarrow \mathbb{R}^+$ such that $\sup_{Q_t^m \in \mathcal{Q}_t^m} \mathbb{E}_{Q_t^m}[S_t^m(X_t)] \leq 1$.
 - Reality announces the (common across games) outcome $x_t \in \mathcal{X}_t$.
 - The skeptic’s wealth in the m -th game gets updated as $\mathcal{K}_t^m := \mathcal{K}_{t-1}^m \cdot S_t^m(x_t)$.

The reader may also refer to Chapter 10 of Shafer and Vovk [30] for related ideas that are presented somewhat differently there.

We now reinterpret our main theorem in a game-theoretic language, inspired by the original results of Ville’s PhD thesis [34], which gave a gambling interpretation to measure-zero sets. In particular, Ville proved that for any event of measure zero (say, under $\text{Ber}(1/2)^\infty$), one can design a betting strategy (a nonnegative martingale) whose wealth increases to infinity whenever that event occurs; see Proposition 8.14 in Shafer and Vovk [29]. Inspired by this result, and the fact that Theorem 7 is a pathwise statement that holds true for every path, we can restate it in the spirit of Ville’s work. To this end, define the set

$$A := \{\omega \in \{0, 1\}^\infty : \text{either the limits in (9) do not exist, or they exist and condition (10) holds}\}.$$

Note that the set A satisfies $\inf_{Q \in \overline{\mathcal{Q}}} Q(A) = 1$, meaning that for an exchangeable sequence, the aforementioned limits (almost) always hold and the corresponding condition is (almost) always satisfied. In other words, its complement satisfies $\sup_{Q \in \overline{\mathcal{Q}}} Q(A^c) = 0$. Then, Theorem 7 states that our process (R_t) increases to infinity whenever A^c occurs. The safety property of (R_t) corresponds to it being a valid betting strategy under the second (and third) game-theoretic protocol presented above. Thus, informally, (R_t) could be seen as an explicitly constructed “witness” to Ville’s theorem in the above context.

Deriving a fully general, ‘robust’ version of Ville’s theorem that holds for composite nulls \mathcal{Q} (and yet recovers Ville’s original theorem for singleton nulls) appears to be an important open direction.

Acknowledgments

AR acknowledges NSF DMS grant 1916320. The authors also thank the reviewers for insightful feedback that led to significant improvements to the paper.

6 References

- [1] Martin Bladt. Characterisation of exchangeable sequences through empirical distributions. *arXiv:1903.07861*, 2019. 27
- [2] D. A. Darling and Herbert Robbins. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences*, 58(1):66–68, July 1967. ISSN 0027-8424, 1091-6490. 23
- [3] A Philip Dawid, Steven de Rooij, Peter Grunwald, Wouter M Koolen, Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Probability-free pricing of adjusted American lookbacks. *arXiv preprint arXiv:1108.4113*, 2011. 25

- [4] A. Philip Dawid, Steven de Rooij, Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Insuring against loss of evidence in game-theoretic probability. *Statistics & Probability Letters*, 81(1):157–162, 2011. ISSN 0167-7152. [25](#)
- [5] Victor H. de la Peña, Michael J. Klass, and Tze Leung Lai. Pseudo-maximization and self-normalized processes. *Probability Surveys*, 4:172–192, 2007. ISSN 1549-5787. [9](#)
- [6] Freddy Delbaen. The structure of m-stable sets and in particular of the set of risk neutral measures. In *In Memoriam Paul-André Meyer: Séminaire de Probabilités XXXIX*, volume 1874 of *Lecture Notes in Math.*, pages 215–258. Springer, Berlin, 2006. [14](#), [33](#)
- [7] Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, 5a edition, 2017. [6](#)
- [8] Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *J. Econom. Theory*, 113(1):1–31, 2003. ISSN 0022-0531. [14](#)
- [9] Hans Föllmer and Alexander Schied. *Stochastic Finance*, volume 27 of *De Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, extended edition, 2004. [32](#)
- [10] Peter Grünwald, Rianne de Heide, and Wouter Koolen. Safe testing. *arXiv:1906.07801*, 2019. [3](#), [6](#), [14](#)
- [11] Edwin Hewitt and Leonard J Savage. Symmetric measures on cartesian products. *Transactions of the American Mathematical Society*, 80(2):470–501, 1955. [26](#)
- [12] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020. [3](#), [6](#), [8](#), [26](#)
- [13] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 2021. [3](#), [9](#), [23](#)
- [14] Garud N. Iyengar. Robust dynamic programming. *Math. Oper. Res.*, 30(2):257–280, 2005. ISSN 0364-765X. [15](#)
- [15] Olav Kallenberg. *Probabilistic Symmetries and Invariance Principles*. Springer Science & Business Media, 2006. [27](#)
- [16] Emilie Kaufmann and Wouter Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *arXiv:1811.11419*, 2018. [6](#), [9](#)
- [17] Werner Kirsch. An elementary proof of de Finetti’s theorem. *arXiv:1809.00882*, 2018. [4](#), [26](#)
- [18] Wouter M. Koolen and Steven de Rooij. Universal codes from switching strategies. *IEEE Transactions on Information Theory*, 59(11):7168–7185, 2013. doi: 10.1109/TIT.2013.2273353. [23](#)
- [19] Wouter M Koolen and Vladimir Vovk. Buy low, sell high. *Theoretical Computer Science*, 558:144–158, 2014. [25](#)
- [20] Raphaël E. Krichevsky and Victor K. Trofimov. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207, March 1981. [10](#), [23](#)
- [21] Tudor Manole and Aaditya Ramdas. Sequential estimation of convex divergences using reverse submartingales and exchangeable filtrations. *arXiv:2103.09267*, 2021. [27](#)
- [22] Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv:2009.03167*, 2020. [3](#), [4](#), [6](#), [7](#), [8](#), [14](#), [23](#), [25](#)
- [23] Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, October 1970. ISSN 0003-4851, 2168-8990. [9](#)

- [24] Herbert Robbins and David Siegmund. Iterated logarithm inequalities and related statistical procedures. In *Mathematics of the Decision Sciences, Part II*, pages 267–279. American Mathematical Society, Providence, 1968. [9](#)
- [25] Herbert Robbins and David Siegmund. Boundary crossing probabilities for the Wiener process and sample sums. *The Annals of Mathematical Statistics*, 41(5):1410–1429, October 1970. ISSN 0003-4851, 2168-8990. [9](#)
- [26] Herbert Robbins and David Siegmund. A class of stopping rules for testing parametric hypotheses. The Regents of the University of California, 1972. [9](#)
- [27] Herbert Robbins and David Siegmund. The expected sample size of some tests of power one. *The Annals of Statistics*, 2(3):415–436, May 1974. ISSN 0090-5364, 2168-8966. [9](#)
- [28] Glenn Shafer. The language of betting as a strategy for statistical and scientific communication (with discussion). *Journal of the Royal Statistical Society, Series A*, 2020. [14](#), [26](#), [27](#)
- [29] Glenn Shafer and Vladimir Vovk. *Probability and finance: it's only a game!*, volume 491. John Wiley & Sons, 2005. [28](#)
- [30] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*, volume 455. John Wiley & Sons, 2019. [8](#), [14](#), [27](#), [28](#)
- [31] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, Bayes factors and p -values. *Statistical Science*, 26(1):84–101, February 2011. ISSN 0883-4237, 2168-8745. [3](#), [6](#), [25](#)
- [32] Alexander Shapiro. Rectangular sets of probability measures. *Operations Research*, 64(2):528–541, 2016. ISSN 0030-364X. [15](#)
- [33] Jun'ichi Takeuchi, Tsutomu Kawabata, and Andrew R. Barron. Properties of Jeffreys mixture for Markov sources. *IEEE Trans. Inf. Theory*, 59(1):438–457, 2013. [10](#), [12](#)
- [34] J Ville. *Étude Critique de la Notion de Collectif*. Gauthier-Villars, Paris, 1939. [28](#)
- [35] Vladimir Vovk. Testing randomness online. *Statistical Science*, 2021. [8](#), [26](#)
- [36] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination, and applications. *The Annals of Statistics*, 2021. [3](#), [25](#)
- [37] Vladimir Vovk, Ivan Petej, Ilia Nouretdinov, Ernst Ahlberg, Lars Carlsson, and Alex Gammerman. Retrain or not retrain: Conformal test martingales for change-point detection. *arXiv:2102.10439*, 2021. [23](#)
- [38] Gordan Žitković. A filtered version of the bipolar theorem of Brannath and Schachermayer. *Journal of Theoretical Probability*, 15(1):41–61, 2002. ISSN 0894-9840. [14](#)
- [39] Abraham Wald. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics*, 16(2):117–186, 1945. [6](#), [8](#)
- [40] Abraham Wald. *Sequential Analysis*. John Wiley & Sons, New York, 1947. [9](#), [14](#)
- [41] Abraham Wald and J Wolfowitz. Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19(3):326–339, 1948. [6](#)
- [42] Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 2020. ISSN 0027-8424. [8](#), [13](#), [23](#), [26](#)
- [43] Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *arXiv:2010.09686*, 2020. [9](#), [14](#), [15](#), [26](#)

- [44] Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013. ISSN 0364-765X. [15](#)
- [45] F. Willems, Y. Shtarkov, and T. Tjalkens. The context-tree weighting method: basic properties. *IEEE Transactions on Information Theory*, 41:653–664, 1995. [12](#)

A Additional technical concepts and definitions

A.1 Reference measures and local absolute continuity

Consider a probability space with a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$. Let \mathbb{R} be a particular probability measure on \mathcal{F}_∞ ; we think of \mathbb{R} as a *reference measure*. We now explain the concept of local domination and how it allows us to unambiguously define conditional expectations.

- If \mathbb{P} is a probability measure on \mathcal{F}_∞ and τ is a stopping time, we write $\mathbb{P}|_\tau$ for the restriction of \mathbb{P} to \mathcal{F}_τ . (This is simply the probability measure on \mathcal{F}_τ defined by $\mathbb{P}|_\tau(A) = \mathbb{P}(A)$, $A \in \mathcal{F}_\tau$. Think of this as the ‘coarsening’ of \mathbb{P} that only operates on events observable up to time τ .)
- \mathbb{P} is called *locally dominated by \mathbb{R}* (or *locally absolutely continuous with respect to \mathbb{R}*), if $\mathbb{P}|_t \ll \mathbb{R}|_t$ for all $t \in \mathbb{N}$. We write this $\mathbb{P} \ll_{\text{loc}} \mathbb{R}$. More explicitly, this means that

$$\mathbb{R}(A) = 0 \quad \Rightarrow \quad \mathbb{P}(A) = 0, \quad \text{for any } A \in \mathcal{F}_t \text{ and } t \in \mathbb{N}.$$

Local absolute continuity does *not* imply that $\mathbb{P} \ll \mathbb{R}$. However, it does imply that $\mathbb{P}|_\tau \ll \mathbb{R}|_\tau$ for any finite (but possibly unbounded) stopping time τ . Indeed, if $A \in \mathcal{F}_\tau$ and $\mathbb{R}(A) = 0$, then $A \cap \{\tau \leq t\} \in \mathcal{F}_t$ for all t , and hence $\mathbb{P}(A) = \lim_{t \rightarrow \infty} \mathbb{P}(A \cap \{\tau \leq t\}) = 0$.

- A set \mathcal{P} of probability measures on \mathcal{F}_∞ is called locally dominated by \mathbb{R} if every element of \mathcal{P} is locally dominated by \mathbb{R} .
- Any $\mathbb{P} \ll_{\text{loc}} \mathbb{R}$ has an associated *likelihood ratio process* (often also called *density process*), namely the \mathbb{R} -martingale (Z_t) given by $Z_t := d\mathbb{P}|_t/d\mathbb{R}|_t$. Being a nonnegative martingale, once Z_t reaches zero it stays there. Thus with the convention $0/0 := 1$, the ratios Z_τ/Z_t are well-defined for any $t \in \mathbb{N}$ and any finite stopping time $\tau \geq t$. Note that each Z_t is defined up to \mathbb{R} -nullsets, and therefore also up to \mathbb{P} -nullsets.
- If $\mathbb{P} \ll_{\text{loc}} \mathbb{R}$ has likelihood ratio process (Z_t) , the following ‘Bayes formula’ holds: for any $t \in \mathbb{N}$, any finite stopping times τ , and any nonnegative \mathcal{F}_τ -measurable random variable Y , one has

$$\mathbb{E}_{\mathbb{P}}[Y \mid \mathcal{F}_t] = \mathbb{E}_{\mathbb{R}} \left[\frac{Z_\tau}{Z_t} Y \mid \mathcal{F}_t \right] \mathbf{1}_{\{Z_t > 0\}}, \quad \mathbb{P}\text{-almost surely.}$$

The right-hand side is uniquely defined \mathbb{R} -almost surely (not just \mathbb{P} -almost surely), and therefore provides a ‘canonical’ version of $\mathbb{E}_{\mathbb{P}}[Y \mid \mathcal{F}_t]$. *We always use this version.* This allows us to view such conditional expectations under \mathbb{P} as being well-defined up to \mathbb{R} -nullsets.

One might ask why we work with *local* domination, rather a ‘global’ condition like $\mathbb{P} \ll \mathbb{R}$ for all \mathbb{P} of interest. The answer is that such a condition would be far too restrictive, as we now illustrate. Let $(X_t)_{t \in \mathbb{N}}$ be a sequence of random variables. For each $\eta \in \mathbb{R}$, let \mathbb{P}^η be the distribution such that the X_t become i.i.d. normal with mean η and unit variance. By the strong law of large numbers, \mathbb{P}^η assigns probability one to the event $A^\eta := \{\lim_{t \rightarrow \infty} t^{-1} \sum_{s=1}^t X_s = \eta\}$. Moreover, the events A^η are mutually disjoint: $A^\eta \cap A^\nu = \emptyset$ whenever $\eta \neq \nu$. This means by definition that the measures \mathbb{P}^μ are all mutually singular. Since there is an uncountable number of them, there cannot exist a measure \mathbb{R} such that $\mathbb{P}^\eta \ll \mathbb{R}$ for all η . On the other hand, if $\mathbb{P}^\eta|_t$ denotes the law of the partial sequence X_1, \dots, X_t for some $t \in \mathbb{N}$, then the measures $\mathbb{P}^\eta|_t$, $\eta \in \mathbb{R}$, are all mutually absolutely continuous. In particular, we could (for instance) use $\mathbb{R} = \mathbb{P}^0$ as reference measure and obtain $\mathbb{P}^\eta \ll_{\text{loc}} \mathbb{R}$ for all $\eta \in \mathbb{R}$.

A.2 Essential supremum

On some probability space, consider a collection $(Y_\alpha)_{\alpha \in \mathcal{A}}$ of random variables, where \mathcal{A} is an arbitrary index set. If \mathcal{A} is uncountable, the pointwise supremum $\sup_{\alpha \in \mathcal{A}} Y_\alpha$ might not be measurable (not a random variable). Moreover, it might happen that $Y_\alpha = 0$ almost surely for every $\alpha \in \mathcal{A}$, but $\sup_{\alpha \in \mathcal{A}} Y_\alpha = 1$. For these reasons, the pointwise supremum is often not useful. Instead, one can use the *essential supremum*.

Proposition 19. *There exists a $[-\infty, \infty]$ -valued random variable Y , called the essential supremum and denoted by $\text{ess sup}_{\alpha \in \mathcal{A}} Y_\alpha$, such that*

1. $Y \geq Y_\alpha$, almost surely, for every $\alpha \in \mathcal{A}$,
2. if Y' is a random variable that satisfies $Y' \geq Y_\alpha$, almost surely, for every $\alpha \in \mathcal{A}$, then $Y' \geq Y$, almost surely.

The essential supremum is almost surely unique.

In words, the essential supremum is the smallest almost sure upper bound on (Y_α) . The proposition guarantees that it always exists. In some cases, more can be said: the essential supremum can be obtained as the limit of an increasing sequence.

Proposition 20. *Suppose (Y_α) is closed under maxima, meaning that for any $\alpha, \beta \in \mathcal{A}$ there is some $\gamma \in \mathcal{A}$ such that $Y_\gamma = \max\{Y_\alpha, Y_\beta\}$. Then there is a sequence (α_n) such that (Y_{α_n}) is an increasing sequence and $\text{ess sup}_{\alpha \in \mathcal{A}} Y_\alpha = \lim_{n \rightarrow \infty} Y_{\alpha_n}$.*

For more information about the essential supremum (and infimum), as well as proofs of the above results, we refer to Section A.5 in [9].

B Omitted proofs

Proof of Proposition 4. It is clear that (ii) implies (i). Now assume that (i) holds, for the moment with $k = 1$. Let M denote the transition matrix of the Markov process (X_t) , i.e., the matrix with elements p_{ji} for $i, j \in \{0, 1\}$. Note that exchangeability yields

$$\mathbb{P}(X_3 = j | X_2 = i) = \mathbb{P}(X_3 = j | X_1 = i)$$

for $i, j \in \{0, 1\}$, hence $M^2 = M$. If M has full rank then M is the identity matrix and (X_t) is constant in time. If M does not have full rank then its two rows are the same since its row sums are always equal to one. It follows that X_1, X_2, \dots are independent. Moreover, by exchangeability X_1, X_2, X_3, \dots have the same distribution. Thus (X_t) is i.i.d. This shows the statement for $k = 1$.

Let us assume that (i) holds for a general $k \in \mathbb{N}$. Instead of providing a modification of the above argument now with transition tensors, let us give an alternative probabilistic, more verbose argument. For simplicity, we only focus on the case $k = 2$; it is clear how to generalize this argument. Note that the Markov property yields that the law of X_5 conditional on (X_1, X_2, X_3, X_4) equals the law of X_5 conditional on (X_3, X_4) . By exchangeability, this also equals the law of X_5 conditional on (X_1, X_2) . This again yields that either (X_t) is constant or the law of X_5 does not depend on the earlier values (X_1, X_2, X_3, X_4) . This observation concludes the proof. \square

Remark 21. *In this remark, let us briefly discuss Proposition 4 in the context of a general Markov process (Y_t) , say with a countable state space. This is not required below but sheds additional light on Proposition 4(ii). To this end, let us hence assume that (Y_t) is an exchangeable Markov process with countable state space. First of all, note that exchangeability yields that (Y_t) reaches state i from state j if and only if it reaches state j from state i . Indeed, for $s, t \in \mathbb{N}$ we have $\mathbb{P}(Y_s = i, Y_t = j) = \mathbb{P}(Y_s = j, Y_t = i)$. Hence, the state space can be partitioned in subsets, say with index set \mathcal{I} , such that in each of these subsets, each state ‘communicates’ with any other. A suitable modification of the proof of Proposition 4 now shows that (Y_t) can be constructed as follows. First draw an \mathcal{I} -valued random variable to choose the subset of the*

state space in which (Y_t) will take values, and then choose i.i.d. draws of a distribution whose support is exactly this subset. As a corollary, conditionally on Y_1 , the random variables Y_2, Y_3, \dots are i.i.d. Indeed, it is easy to see that a $\{0, 1\}$ -valued process (X_t) that satisfies Proposition 4(ii) also satisfies that X_2, X_3, \dots are i.i.d. conditional on X_1 .

Proof of Theorem 14. The proof is essentially a simplified version of an argument due to Delbaen [6, Theorem 11]. This result is argued in continuous time and on a bounded time interval. For the convenience of the reader, we provide a self-contained proof for this paper's discrete-time, infinite-horizon setup.

For each fixed $s \in \mathbb{N}$, L_s is defined as the essential supremum of the family consisting of all $\mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_s]$, indexed by all pairs (\mathbb{Q}, τ) with $\mathbb{Q} \in \mathcal{Q}$ and $\tau \geq s$ a finite stopping time. Here we use a version of the conditional expectation that satisfies $\mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_s] = 0$ on the event $\{Z_s = 0\}$, where (Z_t) denotes the likelihood ratio process of \mathbb{Q} (see Appendix A). We claim that this family of conditional expectations is closed under maxima. To prove this claim, let (\mathbb{Q}, τ) and (\mathbb{Q}', τ') be given. Let $A := \{\mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_s] \geq \mathbb{E}_{\mathbb{Q}'}[E_{\tau'} | \mathcal{F}_s]\}$ and set

$$\tau'' := \tau \mathbf{1}_A + \tau' \mathbf{1}_{A^c}$$

and

$$Z_t'' := \begin{cases} Z_t, & t \leq s \\ \mathbf{1}_A Z_t + \mathbf{1}_{A^c} Z_s \frac{Z_t'}{Z_s'}, & t > s \end{cases}$$

where (Z_t) and (Z_t') are the likelihood ratio processes of \mathbb{Q} and \mathbb{Q}' , respectively. Note that $Z_s' > 0$ on A^c so that Z_t'' is well-defined. Since A belongs to \mathcal{F}_s and $\tau, \tau' \geq s$, τ'' is a (finite) stopping time. Moreover, since \mathcal{Q} is fork-convex, (Z_t'') is the likelihood ratio process of some $\mathbb{Q}'' \in \mathcal{Q}$. We now compute

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}''}[E_{\tau''} | \mathcal{F}_s] &= \mathbb{E}_{\mathbb{Q}''}[\mathbf{1}_A E_\tau | \mathcal{F}_s] + \mathbb{E}_{\mathbb{Q}''}[\mathbf{1}_{A^c} E_{\tau'} | \mathcal{F}_s] \\ &= \mathbb{E}_{\mathbb{R}} \left[\frac{Z_\tau''}{Z_s''} \mathbf{1}_A E_\tau \middle| \mathcal{F}_s \right] + \mathbb{E}_{\mathbb{R}} \left[\frac{Z_{\tau'}''}{Z_s''} \mathbf{1}_{A^c} E_{\tau'} \middle| \mathcal{F}_s \right] \\ &= \mathbb{E}_{\mathbb{R}} \left[\frac{Z_\tau}{Z_s} \mathbf{1}_A E_\tau \middle| \mathcal{F}_s \right] + \mathbb{E}_{\mathbb{R}} \left[\frac{Z_{\tau'}}{Z_s'} \mathbf{1}_{A^c} E_{\tau'} \middle| \mathcal{F}_s \right] \\ &= \mathbf{1}_A \mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_s] + \mathbf{1}_{A^c} \mathbb{E}_{\mathbb{Q}'}[E_{\tau'} | \mathcal{F}_s] \\ &= \max \{ \mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_s], \mathbb{E}_{\mathbb{Q}'}[E_{\tau'} | \mathcal{F}_s] \}. \end{aligned}$$

This demonstrates closure under maxima.

Now fix any $\mathbb{Q} \in \mathcal{Q}$ and $s \in \mathbb{N}$. Thanks to the closure property under maxima, Proposition 20 shows that there exist families (\mathbb{Q}_n) of measures in \mathcal{Q} and (τ_n) of finite stopping times taking values in $\{s, s+1, \dots\}$ such that $\mathbb{E}_{\mathbb{Q}_n}[E_{\tau_n} | \mathcal{F}_s] \uparrow L_s$ almost surely under \mathbb{R} , and hence under \mathbb{Q} . Therefore, by the conditional version of the monotone convergence theorem,

$$\mathbb{E}_{\mathbb{Q}}[L_s | \mathcal{F}_{s-1}] = \mathbb{E}_{\mathbb{Q}} \left[\lim_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Q}_n}[E_{\tau_n} | \mathcal{F}_s] \middle| \mathcal{F}_{s-1} \right] = \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}_n}[E_{\tau_n} | \mathcal{F}_s] | \mathcal{F}_{s-1}]. \quad (17)$$

Replacing \mathbb{Q}_n by $(1 - n^{-1})\mathbb{Q}_n + n^{-1}\mathbb{Q}$ we still have (17) and, in addition, \mathbb{Q} absolutely continuous with respect to \mathbb{Q}_n . From now on we use this modified choice of \mathbb{Q}_n . Let (Z_t) and (Z_t^n) be the likelihood ratio processes of \mathbb{Q} and \mathbb{Q}_n , respectively, and define

$$\tilde{Z}_t^n := \begin{cases} Z_t, & t \leq s, \\ Z_s \frac{Z_t^n}{Z_s^n}, & t > s. \end{cases}$$

By fork-convexity, (\tilde{Z}_t^n) is the likelihood ratio process of some $\tilde{\mathbb{Q}}_n \in \mathcal{Q}$. We then get

$$\begin{aligned}
\mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\tilde{\mathbb{Q}}_n}[E_{\tau_n} | \mathcal{F}_s] | \mathcal{F}_{s-1}] &= \mathbb{E}_{\mathbb{R}} \left[\frac{Z_s}{Z_{s-1}} \mathbb{E}_{\mathbb{R}} \left[\frac{Z_{\tau_n}^n}{Z_s^n} E_{\tau_n} \middle| \mathcal{F}_s \right] \middle| \mathcal{F}_{s-1} \right] \\
&= \mathbb{E}_{\mathbb{R}} \left[\frac{Z_s}{Z_{s-1}} \frac{Z_{\tau_n}^n}{Z_s^n} E_{\tau_n} \middle| \mathcal{F}_{s-1} \right] \\
&= \mathbb{E}_{\mathbb{R}} \left[\frac{\tilde{Z}_{\tau_n}^n}{\tilde{Z}_{s-1}^n} E_{\tau_n} \middle| \mathcal{F}_{s-1} \right] \\
&= \mathbb{E}_{\tilde{\mathbb{Q}}_n} [E_{\tau_n} | \mathcal{F}_{s-1}] \\
&\leq L_{s-1}.
\end{aligned}$$

Combining this with (17) gives $\mathbb{E}_{\mathbb{Q}}[L_s | \mathcal{F}_{s-1}] \leq L_{s-1}$. Iterating this inequality and using that \mathcal{F}_0 is trivial yields $\mathbb{E}_{\mathbb{Q}}[L_s] \leq L_0$. In particular, L_s is \mathbb{Q} -integrable. Since (E_t) is a \mathcal{Q} -safe e-process, we have $L_0 = \sup_{\mathbb{Q} \in \mathcal{Q}, \tau \geq 0} \mathbb{E}_{\mathbb{Q}}[E_{\tau}] \leq 1$. Since $\mathbb{Q} \in \mathcal{Q}$ and $s \in \mathbb{N}$ were arbitrary, this proves that (L_t) is a \mathcal{Q} -NSM with $L_0 \leq 1$.

Let (L'_t) be another \mathcal{Q} -NSM that dominates (E_t) . Then for any $\mathbb{Q} \in \mathcal{Q}$, any $t \in \{0, 1, \dots\}$, and any finite stopping time $\tau \geq t$, the optional stopping theorem under \mathbb{Q} gives $L'_t \geq \mathbb{E}_{\mathbb{Q}}[L'_\tau | \mathcal{F}_t] \geq \mathbb{E}_{\mathbb{Q}}[E_\tau | \mathcal{F}_t]$. Therefore $L'_t \geq L_t$ by the definition of essential supremum. \square