

Anytime valid and asymptotically efficient inference driven by predictive recursion

Vaidehi Dixit* and Ryan Martin†

November 5, 2024

Abstract

Distinguishing two candidate models is a fundamental and practically important statistical problem. Error rate control is crucial to the testing logic but, in complex nonparametric settings, can be difficult to achieve, especially when the stopping rule that determines the data collection process is not available. This paper proposes an e-process construction based on the predictive recursion (PR) algorithm originally designed to recursively fit nonparametric mixture models. The resulting Pre-process affords anytime valid inference and is asymptotically efficient in the sense that its growth rate is first-order optimal relative to PR’s mixture model.

Keywords and phrases: e-process; mixture model; nonparametric; test martingale; universal inference.

1 Introduction

A fundamental problem in statistics is that of distinguishing between two classes of candidate models based on observed data. When the two classes of models are *simple*, i.e., there are just two distinct probability distributions being compared, then [Neyman and Pearson \(1933\)](#) settle the question on how to optimally distinguish the two based on the magnitude of the likelihood ratio. Beyond the test’s statistical properties, the *law of likelihood* (e.g., [Edwards 1992](#); [Hacking 1976](#); [Royall 1997](#)) offers principled justification for comparing the two models in this fashion. In applications, however, it is rare for the relevant comparison to be between two simple hypotheses, and if either of the two hypotheses are *composite*, i.e., consisting of more than one probability distribution, then it is no longer clear how to operationalize the law of likelihood.

Two common strategies are available to quantify the “likelihood” of a composite hypothesis: classical likelihood ratios, like those covered by Wilks’s theorem ([Wilks 1938](#)), maximize the likelihood, whereas Bayes factors (e.g., [Bernardo and Smith 1994](#); [Jeffreys 1998](#)) average the likelihood with respect to suitable priors. In either case, judgments are made as above by considering the magnitude of the corresponding (maximum or average) likelihood ratio. For various reasons, however, simply thresholding these likelihood ratios

*Department of Statistics, University of Missouri, vdixit@missouri.edu

†Department of Statistics, North Carolina State University, rgmart13@ncsu.edu

is not fully satisfactory: classical strategies to calibration are based on sampling distribution calculations that assume a specific data-generating process, and so the reliability statements lack robustness to the kind of departures from these assumptions common in practice, e.g., data-peeking. And since neither of these likelihood ratios generally meet the conditions necessary (see below) to achieve a sort of universal calibration independent of the data-generating process, there is a need for new approaches.

To meet this need, there has been a recent surge of interest in testing procedures that are both valid and efficient under optional stopping; see [Ramdas et al. \(2023\)](#) for a survey of these rapidly-moving developments. The basic building blocks are *e-values*, which in some cases lead to *test martingales* ([Shafer et al. 2011](#)) and other kinds of *e-processes*. These also have close connections to betting interpretations of probability ([Shafer 2021](#)) as well as *game-theoretic probability* ([Shafer and Vovk 2019](#)). A practical challenge, however, is that there are so many e-values for data analysts to choose from in applications. Two general and principled strategies for the construction of e-values (which are sometimes e-processes) are *universal inference* ([Wasserman et al. 2020](#)) and *reverse information projection* ([Grünwald et al. 2024](#)). The former approach has different variants, but these can be computationally demanding and/or lacking in statistical efficiency. The latter strategy is purely efficiency-motivated, and that makes it difficult to apply in nonparametric problems. So, there is a need for an e-process construction that is flexible, conceptually and computationally simple, and is able to achieve the optimal first-order growth rate asymptotically. The present paper offers such a construction.

The familiar likelihood ratios fail to offer *anytime valid* inference in the sense of controlling errors independent of the stopping rule, etc. As explained in [Section 2.2](#), it turns out that e-processes themselves are likelihood ratios. Since the null model is determined by the problem itself, an e-process construction reduces to specification of the non-null model likelihood. Intuitively, one wants this non-null model likelihood to match the “true likelihood” as closely as possible. Of course, the data analyst does not know the true likelihood, so his/her choice of non-null likelihood must be sufficiently flexible to adapt to the information in the data concerning the underlying distribution. Mixtures are flexible, so our proposal is to build this non-null likelihood using a nonparametric mixture model. Fitting such a mixture model is computationally demanding using classical methods, but here we employ a fast, recursive algorithm called *predictive recursion*, or PR for short (e.g., [Newton 2002](#); [Newton et al. 1998](#)); see [Section 2.3](#) for a review. With this flexible and computationally efficient strategy to construct a non-null model likelihood, we define a corresponding e-process in [Section 3](#), which we call the *PRE-process*, that offers anytime valid and efficient statistical inference.

Thanks to PR’s naturally recursive model fit, the PRE-process is easily upper-bounded by a genuine likelihood ratio, which is a test martingale under the null and hence has expected value upper-bounded by 1. Then the anytime validity, e.g., hypothesis tests that are valid uniformly over stopping rules, is an immediate consequence; see [Section 3.2](#). With anytime validity guaranteed, we turn our attention to efficiency in [Section 3.3](#). There we establish that the PRE-process asymptotically achieves the first-order optimal growth rate when the null is false.

Two illustrations are presented in [Section 4](#), where we show that the PRE-process’s empirical, finite-sample growth rate closely matches the theoretical growth rate advertised in [Theorem 2](#). Then, in [Section 5](#), we consider two applications. The first is a classical

nonparametric problem of testing for log-concavity in the density, and we investigate the efficiency of our PRE-process based test compared to the recently proposed tests in [Gangrade et al. \(2023\)](#) and [Dunn et al. \(2024\)](#). The second application concerns testing for time-homogeneity in discrete-valued time series data. Some concluding remarks are in Section [6](#) and further discussion and technical details can be found in the supplementary material.

2 Background

2.1 Notation

The data consists of a sequence X_1, X_2, \dots of \mathbb{X} -valued random variables, and we write $X^n = (X_1, \dots, X_n) \in \mathbb{X}^n$, for $n \geq 1$. Let \mathcal{A} denote the σ -algebra of measurable subsets of \mathbb{X}^∞ and write $\mathcal{A}_k = \sigma(X^k)$, for the filtration generated by X^k , $k \geq 1$. An integer-valued random variable N will be called a *stopping time* if $\{N \leq n\} \in \mathcal{A}_n$.

Consider a collection \mathcal{P} of candidate joint distributions for $X^\infty = (X_1, X_2, \dots)$ on \mathcal{A} and let $P = P^\infty$ denote a generic distribution in \mathcal{P} . Write \mathcal{P}_0 , a subset of \mathcal{P} , for the null hypothesis. Throughout, we use upper-case letters for distributions/probability measures and the corresponding lower-case letters for the associated density (with respect to some common dominating measure). For example, if P_0 is a member of \mathcal{P}_0 , then p_0 is its corresponding density.

2.2 E-processes

An *e-value*, E , relative to P_0 , is a non-negative random variable with P_0 -expected value (upper bounded by) 1. With batch data, say, X , and models that have densities, the non-negativity and expected value constraint implies that $x \mapsto q(x) := E(x) p_0(x)$ is a (sub-probability) density too; therefore, $E(x) = q(x)/p_0(x)$ is a likelihood ratio which can be interpreted as a measure of the evidence in data x against P_0 , relative to Q . Consequently, if the observed value of E is large, then one would be inclined to reject the hypothesis P_0 . A Bayes factor is an e-value relative to a single P_0 ; but likelihood ratios of the common form “ $\sup_p p(x)/p_0(x)$ ” are not e-values.

The situation described above is impractically simple. Indeed, it is rare that \mathcal{P}_0 is a singleton and often data does not come in a batch. *Test martingales* ([Shafer et al. 2011](#)) offer a solution to the non-batch data problem. A test martingale is a non-negative martingale (M_n), relative to P_0 , adapted to the filtration (\mathcal{A}_n) with $M_0 \equiv 1$. Test martingales can also be expressed as likelihood ratios, i.e., if Q is absolutely continuous with respect to P_0 , then

$$M_n = \frac{q(X^n)}{p_0(X^n)} = \prod_{i=1}^n \frac{q(X_i | X^{i-1})}{p_0(X_i | X^{i-1})}, \quad n \geq 1. \quad (1)$$

So then the construction of a test martingale boils down to the choice of Q for the numerator. To handle composite \mathcal{P}_0 cases, one can construct a collection of test martingales ($M_n^{P_0}$) indexed by n and by $P_0 \in \mathcal{P}_0$. This also has a likelihood ratio form but, the choice of Q could vary with P_0 . Now define an *e-process* (E_n) as a sequence of non-negative

random variables with $E_n \leq M_n^{P_0}$ for all n and all $P_0 \in \mathcal{P}_0$. An example of this would be $E_n = \inf_{P_0 \in \mathcal{P}_0} M_n^{P_0}$ and, if it happened that the Q in (1) does not depend on P_0 , then this simplifies to

$$E_n = \frac{q(X^n)}{\sup_{P_0 \in \mathcal{P}_0} p_0(X^n)}, \quad n \geq 1.$$

This is a common e-process construction (e.g., Wasserman et al. 2020, Sec. 8) but it is not the only one; see Ramdas et al. (2023) for more discussion on the various alternatives, in particular, the proposals in Grünwald et al. (2024) and in Waudby-Smith and Ramdas (2024). We focus here on e-processes that, as above, only depend on a choice of Q representing the alternative.

E-processes tend to be small under \mathcal{P}_0 . In particular, *Ville's inequality* (e.g., Howard et al. 2021; Shafer and Vovk 2019) states that if M_n is a test martingale for \mathcal{P}_0 , then

$$\sup_{P_0 \in \mathcal{P}_0} P_0(M_N \geq \alpha^{-1}) \leq \alpha, \quad \text{for all } \alpha \in (0, 1) \text{ and all stopping times } N.$$

Of course, the same is true for any e-process E_n bounded by M_n . The implications of this are far-reaching: it leads to statistical tests that are *anytime valid* in the sense that the reliability claims hold (basically) no matter how the investigator decides to conclude their study and perform the statistical analysis, thereby offering additional “safety” (cf., Grünwald et al. 2024) compared to procedures that only control the Type I error rate for a particular sampling scheme. For example, a test that rejects \mathcal{P}_0 based on data X^n when $E_n \geq \alpha^{-1}$ controls the Type I error rate *even if the investigator peeked at the data* when deciding whether to conclude the study at time n .

2.3 Predictive recursion

Here and in what follows, we assume X_1, X_2, \dots are iid. Following the notation above, a general mixture model for the (common) marginal distribution defines the density function as

$$q^\Psi(x) = \int_{\mathbb{U}} p_u(x) \Psi(du), \quad x \in \mathbb{X}, \quad (2)$$

where p_u is the density corresponding to a distribution P_u , with $\{P_u : u \in \mathbb{U}\} \subseteq \mathcal{P}$, and Ψ is a mixing distribution defined on \mathbb{U} . Such models are incredibly flexible, so they are often used in contexts where the mixture structure itself is not directly relevant, e.g., in density estimation applications. It is common to think of $x \mapsto p_u(x)$ as a parametric kernel, like Gaussian, but that is not necessary here; indeed, \mathbb{U} could just be a generic indexing of the entire model \mathcal{P} . But we should emphasize that mixtures of simple parametric kernels can be incredibly flexible. For example, the set of Gaussian location–scale mixtures is dense with respect to total variation distance in the space of smooth densities in Euclidean space (e.g. DasGupta 2008, Theorem 33.1).

For fitting the mixture model (2) to the observed data $X^n = (X_1, \dots, X_n)$, one common strategy is nonparametric maximum likelihood (e.g., Laird 1978; Lindsay 1995). Another common strategy is to introduce a prior distribution for Ψ , e.g., a Dirichlet process prior (e.g., Ferguson 1973; Ghosal 2010; Lo 1984), and carry out a nonparametric Bayesian analysis. There are advantages to sticking with existing approaches, the

complexity of these models creates computational challenges. In particular, since tailored Monte Carlo methods are required, one cannot capitalize on the Bayesian coherent updating property—“today’s posterior is tomorrow’s prior”—when data are processed sequentially. Maximum likelihood estimation faces similar challenges. A novel alternative, developed by Michael Newton and collaborators in the late 1990s, is a fast, recursive algorithm called *predictive recursion*, or PR for short (e.g., [Newton 2002](#)). Applications in large-scale inference settings include [Tao et al. \(1999\)](#), [Newton et al. \(2001\)](#), [Martin and Tokdar \(2012\)](#), [Tansey et al. \(2018\)](#), and [Woody et al. \(2022\)](#).

Here we provide a quick review of PR and its properties; for more details, see [Martin \(2021b\)](#). Start with a weight sequence $(w_i : i \geq 1) \subset (0, 1)$ that satisfies

$$\sum_{i=1}^{\infty} w_i = \infty \quad \text{and} \quad \sum_{i=1}^{\infty} w_i^2 < \infty. \quad (3)$$

Following [Martin and Tokdar \(2009\)](#) and the discussion at the end of Section [3.3](#) below, we take $w_i = (i + 1)^{-0.67}$ for all of our examples. Next, take an initial guess $\hat{\Psi}_0$ of Ψ supported on the index space \mathbb{U} . Then PR updates the initial guess along the data sequence as follows:

$$\hat{\Psi}_i(du) = (1 - w_i) \hat{\Psi}_{i-1}(du) + w_i \frac{p_u(X_i) \hat{\Psi}_{i-1}(du)}{\int_{\mathbb{U}} p_v(X_i) \hat{\Psi}_{i-1}(dv)}, \quad u \in \mathbb{U}, \quad i \geq 1. \quad (4)$$

PR is recursive, so one only needs $\hat{\Psi}_n$ and the new data point X_{n+1} to get the new $\hat{\Psi}_{n+1}$. Moreover, this leads naturally to a plug-in estimator of the mixture density

$$\hat{q}_{x^n}(x) := q^{\hat{\Psi}_n}(x) = \int_{\mathbb{U}} p_u(x) \hat{\Psi}_n(du), \quad x \in \mathbb{X}.$$

Large-sample properties of PR have been explored in [Ghosh and Tokdar \(2006\)](#), [Martin and Ghosh \(2008\)](#), [Tokdar et al. \(2009\)](#), [Martin and Tokdar \(2009, 2011\)](#), and [Dixit and Martin \(2023b\)](#). The property most relevant to our efforts here will be explained in Section [3](#).

The above display defines a PR-based predictive density for X_{n+1} , given X^n . It also produces a joint marginal density for X^n —“marginal” in the sense that the mixing distribution Ψ has been integrated out—and it has a multiplicative form, à la [Wald \(1947, Eq. 10.10\)](#), [Dawid \(1984\)](#), and [Rissanen \(1984\)](#). That is, the joint marginal density \hat{q}^{PR} satisfies

$$\hat{q}^{\text{PR}}(x^n) = \hat{q}_{x^{n-1}}(x_n) \hat{q}^{\text{PR}}(x^{n-1}) = \prod_{i=1}^n \hat{q}_{x^{i-1}}(x_i), \quad n \geq 1, \quad x^n \in \mathbb{X}^n. \quad (5)$$

It is PR’s flexibility and the multiplicative form of its joint marginal density that makes it especially suitable for anytime valid nonparametric inference.

The simple PR algorithm is computationally efficient in the sense that the update $\hat{\Psi}_n \rightarrow \hat{\Psi}_{n+1}$ is an $O(1)$ operation in n . But the integration in [\(4\)](#) must be done numerically and requires care. If the dimension of \mathbb{U} is relatively low, then integration can be handled easily and accurately using quadrature. Monte Carlo would be a natural alternative,

but direct sampling from $\hat{\Psi}_n$ is a challenge, so we developed a “PRticle filter” strategy (Dixit and Martin 2023a) that can easily accommodate mixtures over \mathbb{U} of moderate dimension, as in Example 2 below and in the supplementary material. Extensions to high-dimensional \mathbb{U} is a focus of ongoing research.

The computations in (4) are carried out pointwise on a grid spanning \mathbb{U} , which effectively requires \mathbb{U} to be fixed *a priori* and bounded. In our experience, particularly with batch data, one can safely fix \mathbb{U} as a large compact that contains those u ’s even remotely compatible with the observed data. With sequential data, one could first use a block of data to inform the choice of \mathbb{U} , then go back and process the full sequence with that choice of \mathbb{U} . More formally, (Martin and Tokdar 2009, Sec. 5) present a *generalized predictive recursion* algorithm that features an adaptive support, but we will not consider this here.

3 PRe-processes

3.1 Construction

Consider null hypothesis $H_0 : P \in \mathcal{P}_0$. The proposal here is to construct a suitable marginal likelihood under \mathcal{P}_0^c by mixing over a specified class $\{P_u : u \in \mathbb{U}\}$ of distributions. Instead of a computationally demanding nonparametric maximum likelihood or Bayes fit, here we make use of the efficient PR algorithm reviewed in Section 2.3 above.

Let Ψ be a mixing probability distribution supported on the specified index set \mathbb{U} and consider a basic mixture model as in (2). This model can be fit to data X^n using PR, and the output most relevant to us here is the joint marginal density for X^n that is produced as a by-product, namely, $\hat{q}^{\text{PR}}(X^n)$ as given in (5). Recall the multiplicative form of $\hat{q}^{\text{PR}}(X^n)$ in (5), a driving force behind the general formulation in (Wasserman et al. 2020, Sec. 8). Moreover, the PR update from $\hat{q}^{\text{PR}}(X^{n-1})$ to $\hat{q}^{\text{PR}}(X^n)$ is an $O(1)$ computation, compared to the $O(n)$ computation expected with “non-anticipatory” maximum likelihood estimation (e.g., (Gangrade et al. 2023)).

With this, we define the following PR-driven e-process, i.e., *PRE-process*,

$$E_n^{\text{PR}} = E^{\text{PR}}(X^n; \mathcal{P}_0) = \frac{\hat{q}^{\text{PR}}(X^n)}{\sup_{p \in \mathcal{P}_0} p(X^n)}, \quad n \geq 1. \quad (6)$$

Intuitively, since the likelihood tends to favor the true hypothesis, if H_0 is true (resp. false), then E_n^{PR} ought to be small (resp. large). This intuition suggests a test that rejects H_0 if and only if E_n^{PR} is large, and we justify this intuition in Section 3.2 below.

Our proposal can be compared to two of the now-standard e-process constructions reviewed in (Ramdas et al. 2023). Specifically, our proposal is, on the one hand, like the basic *method of mixtures* (e.g., (Darling and Robbins 1968; Wald 1945)) in that it forms a marginal likelihood under the alternative via a suitable mixture model, which we think is intuitively appealing. To achieve both the flexibility and the appealing intuition, the mixture model needs to be nonparametric, which would pose computational challenges for traditional likelihood-based methods. But PR is specifically designed to simply and efficiently handle this challenge. So, the PRE-process is like (the no data-splitting variants of) universal inference (Wasserman et al. 2020, Sec. 8), just with a specific focus on flexible mixture model alternatives with a fast, recursive updating scheme.

Our proposal can also be compared to that in Grünwald et al. (2024). While our choice to use “ $\sup_{p \in \mathcal{P}_0} p(X^n)$ ” in the denominator of (6) is a natural one, it is not the only option. Grünwald et al. propose a strategy based on the *reverse information projection* (RIP), as in Li and Barron (2000), which amounts to replacing “ $\sup_{p \in \mathcal{P}_0} p(X^n)$ ” in the denominator with the likelihood at a fixed—but strategically chosen—density, say, p_0^{RIP} in the convex hull, $\text{co}(\mathcal{P}_0)$, of \mathcal{P}_0 . In our context, they might propose to make inference based on the ratio

$$E_n^{\text{PR+RIP}} = \hat{q}^{\text{PR}}(X^n) / p_0^{\text{RIP}}(X^n),$$

where P_0^{RIP} satisfies $K(\hat{Q}^{\text{PR}}, P_0^{\text{RIP}}) = \inf_{P_0 \in \text{co}(\mathcal{P}_0)} K(\hat{Q}^{\text{PR}}, P_0)$. This strategy has powerful motivation and nice properties concerning maximal growth rate. However, the complexity of the applications we have in mind, along with PR’s non-trivial learning process, raise some difficult questions: first, how to compute P_0^{RIP} and, second, does $E_n^{\text{PR+RIP}}$ define a proper e-process? We leave these questions for future investigation.

3.2 Validity

Of course, we cannot refer to the quantity defined in (6) as an “e-process” without showing that it satisfies the required properties. Theorem 1 establishes the basic e-process property from which all the relevant statistical properties follow.

Theorem 1. *Consider a model \mathcal{P} for the iid data sequence X_1, X_2, \dots . For a model $\mathcal{P}_0 \subset \mathcal{P}$ to be tested, the PRE-process defined in (6) is an e-process.*

Proof. For any fixed n , since the supremum over \mathcal{P}_0 appears in the denominator of (6), the following inequality is immediate:

$$E_n^{\text{PR}} = E^{\text{PR}}(X^n; \mathcal{P}_0) \leq E^{\text{PR}}(X^n; \{p_0\}) = \hat{q}^{\text{PR}}(X^n) / p_0(X^n), \quad \text{for all } P_0 \in \mathcal{P}_0.$$

The upper bound is a collection of test martingales indexed by $P_0 \in \mathcal{P}_0$ and, therefore, (E_n^{PR}) is an e-process under \mathcal{P}_0 . \square

From here, we can immediately deduce several directly interpretable statistical results. In particular, suitably-defined PRE-process-based testing procedures are anytime valid.

Corollary 1. *For a desired significance level $\alpha \in [0, 1]$, let $T_\alpha(X^n)$ be the test that rejects $H_0 : P \in \mathcal{P}_0$ if and only if $E^{\text{PR}}(X^n; \mathcal{P}_0) \geq \alpha^{-1}$. This test controls the frequentist Type I error at the designated level, i.e., $\sup_{P_0 \in \mathcal{P}_0} P_0\{T_\alpha(X^N) \text{ rejects}\} \leq \alpha$ for any stopping rule N .*

If one has a valid test, then of course it can be inverted to construct valid confidence sets. Beyond these familiar statistical procedures, one can leverage the PRE-process for the purpose of broader uncertainty quantification, as discussed in, e.g., Martin (2022, 2023b) and Grünwald (2022, 2023). For further details on these, see the supplementary material.

3.3 Asymptotic growth rate

Theorem 1 establishes that PRE-process-based tests, etc. are anytime valid. For the PRE-process procedure to be *efficient*, we want the e-value to be large under the alternative, at least asymptotically, so that we will correctly reject false null hypotheses. Our goal here is to show that the PRE-process is asymptotically optimal in the sense of Cover and Thomas (2006, Ch. 11), i.e., its asymptotic growth rate is optimal to first order in the exponent. Towards this, consider the case where $P^* \notin \mathcal{P}_0$ determines the true distribution of the data X^∞ . So far, we have been implicitly assuming that mixtures of the kernels $\{P_u : u \in \mathbb{U}\}$ are sufficiently flexible to warrant consideration as a “non-null model,” i.e., one is willing to use such a model for nonparametric estimation of the true distribution; see Section 2.3. If \mathcal{Q} denotes the convex hull of $\{P_u : u \in \mathbb{U}\}$, the set of all mixtures of these kernels, then a more explicit statement of the aforementioned implicit assumption is that \mathcal{Q} is an acceptable non-null model. In other words, if one believes that $P^* \notin \mathcal{P}_0$, then he/she would also believe that P^* is “closer” to \mathcal{Q} than to \mathcal{P}_0 .

The main result in Martin and Tokdar (2011) states that, under certain regularity conditions (see the supplementary material), with P^* -probability 1,

$$n^{-1} \log \hat{q}^{\text{PR}}(X^n) = n^{-1} \log p^*(X^n) - K(P^*, \mathcal{Q}) + o(1), \quad n \rightarrow \infty, \quad (7)$$

where $K(P^*, \mathcal{Q})$ is the Kullback–Leibler divergence of \mathcal{Q} from P^* and

$$K(P^*, \mathcal{Q}) = \inf_{Q \in \mathcal{Q}} K(P^*, Q), \quad (8)$$

with $\mathcal{Q} = \text{co}(\{P_u : u \in \mathbb{U}\})$ the set of mixtures. For the null \mathcal{P}_0 , let $\kappa^*(\mathcal{P}_0)$ be such that

$$\liminf_{n \rightarrow \infty} n^{-1} \log \{p^*(X^n)/\hat{p}_0(X^n)\} \geq \kappa^*(\mathcal{P}_0), \quad \text{with } P^*\text{-probability 1,} \quad (9)$$

where $\hat{p}_0(X^n) = \sup_{P_0 \in \mathcal{P}_0} p_0(X^n)$ is the maximum likelihood estimator under the null \mathcal{P}_0 . From the definition of \hat{p}_0 , it follows that $\kappa^*(\mathcal{P}_0) \leq K(P^*, \mathcal{P}_0) := \inf_{P_0 \in \mathcal{P}_0} K(P^*, P_0)$, and often $\kappa^*(\mathcal{P}_0)$ will equal this minimal Kullback–Leibler number. For example, if $\mathcal{P}_0 = \{P_0\}$ is a singleton and $K(P^*, P_0) < \infty$, then the law of large numbers gives (9) with “lim,” “=,” and $\kappa^*(\mathcal{P}_0) = K(P^*, P_0)$. More generally, suppose there exists $P_0^\dagger \in \mathcal{P}_0$ with $K(P^*, P_0^\dagger) = K(P^*, \mathcal{P}_0)$; see, e.g., Patilea (2001) and Kleijn and van der Vaart (2006). Then

$$n^{-1} \log \{p^*(X^n)/\hat{p}_0(X^n)\} = n^{-1} \log \{p^*(X^n)/p_0^\dagger(X^n)\} + n^{-1} \log \{p_0^\dagger(X^n)/\hat{p}_0(X^n)\}.$$

The first term on the right-hand side converges to $K(P^*, \mathcal{P}_0)$ with P^* -probability 1, so, the bound in (9) is determined by the second term on the right-hand side. That term vanishes under relatively mild conditions on \mathcal{P}_0 , e.g., when the maximum likelihood estimator is consistent for P_0^\dagger (relative to the true distribution P^*). In these typical cases, (9) holds with “lim,” “=,” and $\kappa^*(\mathcal{P}) = K(P^*, \mathcal{P}_0)$. But it is possible that $\kappa^*(\mathcal{P}_0) < K(P^*, \mathcal{P}_0)$, e.g., in cases where consistency fails; see the supplementary material for an example.

Putting all this together, with P^* -probability 1, the PRE-process (6) satisfies

$$\begin{aligned} n^{-1} \log E_n^{\text{PR}} &= n^{-1} \log \{\hat{q}^{\text{PR}}(X^n)/\hat{p}_0(X^n)\} \\ &= n^{-1} \log \{p^*(X^n)/\hat{p}_0(X^n)\} + n^{-1} \log \{\hat{q}^{\text{PR}}(X^n)/p^*(X^n)\} \\ &\geq \kappa^*(\mathcal{P}_0) - K(P^*, \mathcal{Q}) + o(1). \end{aligned}$$

More details about (7), (9), and the proof are given in the supplementary material.

Theorem 2. Consider X_1, X_2, \dots iid with distribution P^* . If (7) and (9) hold, then with P^* -probability 1, the PRE-process has asymptotic first-order growth rate

$$\Delta(P^*; \mathcal{P}_0, \mathcal{Q}) = \kappa^*(\mathcal{P}_0) - K(P^*, \mathcal{Q}). \quad (10)$$

That is, as $n \rightarrow \infty$, the PRE-process satisfies

$$\log E_n^{\text{PR}} \geq n \Delta(P^*; \mathcal{P}_0, \mathcal{Q}) + o(n), \quad \text{with } P^*\text{-probability 1.}$$

A few remarks about Theorem 2 are in order. We focus our attention here on the typical case where $\kappa^*(\mathcal{P}_0) = K(P^*, \mathcal{P}_0)$. First, the non-null case where $K(P^*, \mathcal{P}_0) > K(P^*, \mathcal{Q})$ is of primary interest. This, again, is because \mathcal{Q} is judged to be an acceptable nonparametric model, so surely a non-null P^* is closer to \mathcal{Q} than to \mathcal{P}_0 . We claim that the growth rate $\Delta(P^*; \mathcal{P}_0, \mathcal{Q})$ is “optimal” which deserves explanation. In the simple-versus-simple case, with $\mathcal{P}_0 = \{P_0\}$ and $\mathcal{Q} = \{P_1\}$, “the likelihood ratio is growth rate optimal” (Ramdas et al. 2023) and its growth rate $K(P_1, P_0)$ agrees with (10). More generally, given a commitment to maximum \mathcal{P}_0 -likelihood in the e-process denominator, the user’s only input in the e-process construction is the choice of “ $\hat{q}(X^n)$ ” in the numerator. If this \hat{q} is constrained to \mathcal{Q} , then there is no choice of numerator that can achieve a larger first-order asymptotic growth rate than that in (10). This can also be compared to the growth rates in Grünwald et al. (2024). Indeed, if the true P^* was known and used in their e-value’s numerator, then their solution is, by definition, growth rate optimal and its first-order growth rate agrees with that in Theorem 2. Since the growth rate cannot be better when the true P^* is unknown, our PRE-process must be asymptotically optimal in the above sense too.

Second, suppose $K(P^*, \mathcal{P}_0) < K(P^*, \mathcal{Q})$, i.e., \mathcal{P}_0 is “more true” than \mathcal{Q} ; a special case is when $K(P^*, \mathcal{P}_0) = 0$ and hence \mathcal{P}_0 is “true.” Then we want the PRE-process to be vanishing with n and, in the not-uncommon case where (9) holds with “lim” and “=,” Theorem 2 establishes this: if $K(P^*, \mathcal{P}_0) < K(P^*, \mathcal{Q})$, then the PRE-process vanishes as $n \rightarrow \infty$; see, also, Ramdas et al. (2022, Sec. 8.2). In the edge case where $K(P^*, \mathcal{P}_0) = K(P^*, \mathcal{Q})$, so \mathcal{P}_0 and \mathcal{Q} are “equally true” and cannot be distinguished, we do not specifically expect E_n^{PR} to grow or to vanish, so Theorem 2’s ambiguity about the behavior of E_n^{PR} in such cases accurately reflects this.

Next, how small is the $o(n)$ term in Theorem 2’s lower bound? This is relevant because it describes the asymptotic gap between the proposed PRE-process and an oracle e-process that uses the true p^* in the numerator. Specifically, the oracle e-process is given by

$$E_n^{\text{OR}} = \frac{p^*(X^n)}{\sup_{P_0 \in \mathcal{P}_0} p_0(X^n)},$$

and, asymptotically, the difference between the two log-e-processes, $\log E_n^{\text{OR}} - \log E_n^{\text{PR}}$, is the “ $o(n)$ ” term in our lower bound. Moreover, that same log-difference is

$$\log E_n^{\text{OR}} - \log E_n^{\text{PR}} = \log \frac{p^*(X^n)}{\hat{q}^{\text{PR}}(X^n)} = \sum_{i=1}^n \log \frac{p^*(X_i)}{\hat{q}_{X^{i-1}}(X_i)},$$

where $\hat{q}_{X^{i-1}}$ is PR’s one-step-ahead predictive density as in (5). Martin and Tokdar (2011) showed that the right-hand side of the above display grows, asymptotically, like

$nK(p^*, \hat{q}_{X^n})$. Furthermore, [Martin and Tokdar \(2009\)](#) proved the following: in addition to the conditions presented in the supplementary material, if P^* is in the interior of \mathcal{Q} and if PR’s weight sequence (w_i) is such that $w_n \asymp n^{-\gamma}$ for $\gamma \in (2/3, 1]$, then $n^{1-\gamma}K(p^*, \hat{q}_{X^n}) \rightarrow 0$ with P^* -probability 1. One immediate take-away is that, apparently, $\gamma \approx 2/3$ gives the best rates for PR. Those authors argue that their rate is “worst-case” corresponding to P^* arbitrarily close to the boundary of \mathcal{Q} , which is most difficult for PR to learn. When P^* is away from the boundary of \mathcal{Q} , however, PR’s rate tends to be faster than Martin & Tokdar’s conservative upper bound. In fact, our empirical work suggest the following conjecture: $n^{1-\gamma/2}K(p^*, \hat{q}_{X^n}) \rightarrow 0$; see [Section 6](#). If true, then this would be near-optimal in certain cases. For instance, if p^* is a monotone decreasing density, i.e., a scale mixture of uniform kernels (e.g., [Williamson 1956](#)), then PR’s rate with $\gamma \approx 2/3$ would virtually agree with the well-known minimax optimal rate of $n^{-2/3}$ in Kullback–Leibler divergence. (There is another new sense in which PR’s rate is actually faster than what Martin & Tokdar’s result suggests, which we present in [Section 4](#) of the supplementary material.) Returning to the original question, the aforementioned theory implies that, with $\gamma \approx 2/3$, the PRE-process’s $o(n)$ error term is conservatively like $n^{2/3}$. Our conjecture, however, says that the error is more like $n^{1/3}$, and the empirical results in [Section 5](#) of the supplementary material support this. Better approximations of the oracle e-process are possible in some cases, e.g., using nonparametric Bayes predictive densities in the e-process numerator, but not with PR’s computational efficiency—and the corresponding “ $o(n)$ ” term would still generally be of polynomial order in n .

4 Illustrations

Here we consider two illustrations of the procedure described in [Section 3](#): testing for monotonicity and testing a parametric model. Our focus here is on comparing the empirical growth rate of the PRE-process to that predicted by [Theorem 2](#). So, in what follows, the data X_1, X_2, \dots will be generated from a distribution P^* that does not belong to \mathcal{P}_0 .

Example 1. Monotone densities on the positive half-line are common in biomedical, engineering, and astronomy applications. Our goal is to test the null hypothesis that the underlying density is monotone. For the denominator of our PRE-process, we use R package [REBayes \(Koenker and Gu 2017\)](#) to find the nonparametric maximum likelihood estimator, i.e., the Grenander estimator (e.g., [Grenander 1956](#)). For the numerator, we use a mixture model with a gamma kernel, i.e., p_u is a gamma density with u the shape and rate parameter pair. Then we fit a mixture model over $\mathbb{U} = [1, 15] \times [10^{-5}, 5]$ with the initial guess Ψ_0 the uniform distribution over \mathbb{U} . For our experiments, the true distribution P^* is a gamma distribution with unit rate parameter and varying shape parameter. The idea is that, if the shape parameter is 1, then P^* would be exponential which is monotone; so as the shape parameter varies from 2, to 5, and to 10, the density becomes “less monotone.” We generate 100 data sets from P^* and calculate the PRE-process E_n^{PR} at the increments $n = 100, 200, \dots, 5000$. A plot of the $\log E_n^{\text{PR}}$ versus n is displayed in [Figure 1\(a\)](#), along with a reference line having slope $K(P^*, \mathcal{P}_0)$. First note that, as expected, the growth rate increases as the true density gets “less monotone” and, second, that the PRE-process closely follows the theoretical growth rate across all three scenarios.




Figure 1: Plots of $\log E_n^{\text{PR}}$ versus n for 100 data sets of size $n = 5000$ under three settings of the true density in the alternative. Each case shows a line with slope based on the theoretical growth rate in Theorem 2. In both panels, the gray, blue, and pink lines correspond to cases where the data are generated from a distribution that is “close,” “less close,” and “farthest” from the null, respectively. Specific details about the settings are given in the text.

Example 2. Testing for clustering of directional data is an important practical problem (e.g., Mardia and Jupp 2000). Here we focus on testing for a single von Mises–Fisher model \mathcal{P}_0 versus a mixture thereof. This problem is also challenging since the null lies on the boundary of the alternative (e.g., McVinish et al. 2009; Tokdar et al. 2010).

Computation of the PRE-process’s denominator is straightforward, since closed-form expressions are available for the maximum likelihood estimators. The von Mises–Fisher distribution is generally parameterized with a mean parameter μ and concentration parameter κ , where $\|\mu\| \leq 1$ and $\kappa \geq 0$. So for the numerator we take the mixing distribution variable as $u = (\mu_\theta, \mu_\phi, \kappa)$ where $(\mu_\theta, \mu_\phi) \in (0, \pi) \times (0, 2\pi)$ represent the spherical coordinates of μ . Note that because of the multivariate nature of the problem we need to implement the PRticle filter machinery in Dixit and Martin (2023a) to calculate the normalizing constants. For this we generate particles of size $t = 10,000$ from $\text{Unif}(0, \pi)$, $\text{Unif}(0, 2\pi)$ and $\text{Unif}(0.2, 10)$ respectively for μ_θ , μ_ϕ and κ . For our simulation, we generate data from a bimodal mixing distribution Ψ^* which fixes $\kappa = 10$ and distributes (μ_θ, μ_ϕ) over $(0, \pi) \times (0, 2\pi)$ according to a mixture of truncated bivariate normals,

$$(\mu_\theta, \mu_\phi) \sim 0.5 \text{trN}_2(\eta_1, \Sigma) + 0.5 \text{trN}_2(\eta_2, \Sigma). \quad (11)$$

In (11), the covariance matrix Σ is a diagonal matrix with the diagonal vector as $\{(\pi/12)^2, (\pi/6)^2\}$, and the mean vectors are $\eta_1 = (\pi/4, \pi/2)^\top$ and η_2 varying,

$$\eta_2 \in \{(3\pi/8, 3\pi/4)^\top, (\pi/2, \pi)^\top, (3\pi/4, 3\pi/2)^\top\}$$

. The distance between the two modes (η_1, η_2) acts as a measure of the “degree of clustering,” with large distances corresponding to going *further away* from the null. We generate 100 data sets under each configuration and plot $\log E_n^{\text{PR}}$ versus n at increments $n = 100, 200, \dots, 5000$ in Figure 1(b), along with a reference line of slope $K(P^*, \mathcal{P}_0)$. As expected, the slope increases in the “degree of clustering” and the log-PRE-process paths closely follow this trend.

5 Applications

5.1 Testing log-concavity of a density

Examples of log-concave densities on the real line include the Gaussian, logistic, Laplace, and others. Efficient numerical methods have been developed (e.g., [Cule et al. \(2010\)](#); [Dümbgen and Rufibach \(2011\)](#)) to find the maximum likelihood estimator of a log-concave density, and its asymptotic properties are studied in, e.g., [Doss and Wellner \(2016\)](#). As for testing log-concavity, e-process-based tests have been proposed recently in [Gangrade et al. \(2023\)](#) and [Dunn et al. \(2024\)](#). Here, we compare the PRE-process’s growth rate to that of the aforementioned e-processes.

For the PRE-process denominator, we get the log-concave maximum likelihood estimator using the `logConDens` function in the R package `logcondens` ([Dümbgen and Rufibach \(2011\)](#)). For the numerator, we consider a mixture model with Gaussian kernel p_u , where u is mean and standard deviation pair, and mixing distribution supported on $\mathbb{U} = [-10, 20] \times [0.01, 3]$. An R package for PRE-process testing of log-concavity can be found at <https://github.com/vdixit005/PReprocess>. For Dunn et al.’s e-process, E_n^{UI} , we use the code at <https://github.com/RobinMDunn/LogConcaveUniv> with its default settings. For the e-process proposed in [Gangrade et al. \(2023\)](#), E_n^{ULR} , we follow their recommendation and use an iteratively fit Gaussian kernel density estimator.

For the simulation, we take P^* as an uneven mixture of two normals, with means 0 and μ and variances 2. As above, μ also acts as a measure of the “degree of non-log-concavity,” i.e., if μ is close to 0, then p^* is only mildly non-log-concave; otherwise, p^* is more severely non-log-concave. The three cases we consider in our experiments are $\mu = 6, 10, 14$. We generate 100 data sets under each, calculate E_n^{PR} , E_n^{UI} and E_n^{ULR} at increments $n = 100, 200, \dots, 5000$. Plots of $\log E_n$ versus n , along with a pointwise average over replications for each method are shown in [Figure 2](#). The key takeaways are as follows. First, in terms of statistical efficiency, our PRE-process’s growth rate is faster than Dunn et al.’s and no slower than Gangrade et al.’s in each scenario. Interestingly, the slopes of the n versus $\log E_n^{\text{UI}}$ lines in [Figure 2](#) are roughly half that of the n versus $\log E_n^{\text{PR}}$ lines, which suggests that the loss of efficiency is due to the 50–50 data-splitting Dunn et al. employ. Second, in terms of stability, the log-PRE-process has considerably smaller variance than Gangrade et al.’s log-e-process. Third, in terms of computational efficiency, both E_n^{UI} and E_n^{ULR} process the entire observed data sequence each time a batch of data arrives because the previous calculation cannot be directly updated. The PRE-process, on the other hand, has a genuinely recursive update and is much faster.

The above comparison focused on the scalar-data case, but the same e-processes can be defined and used for testing log-concavity in multivariate settings. In the supplementary material we present a comparison of E_n^{PR} and E_n^{UI} in a multivariate case. The take-away message there is that increasing the dimension naturally impacts the quality of the tests, but it affects the PRE-process less in the sense that E_n^{PR} tends to be closer to a suitable oracle e-process than E_n^{UI} does.




Figure 2: Plots of $\log E_n^{\text{PR}}$ (blue), $\log E_n^{\text{ULR}}$ (pink) and $\log E_n^{\text{UI}}$ (green) versus n , based on the not-log-concave density p^* as a mixture of normals. Panel (a) corresponds to p^* “closest” to log-concave while (c) is “farthest.”

5.2 Testing time homogeneity

Let $X_i = (X_{i,1}, \dots, X_{i,T})$ denote a binary time series of length $T \geq 1$, i.e., $X_{i,t} \in \{0, 1\}$ for each (i, t) pair, generally dependent within the X_i ’s, but independent between X_i and X_j , for $j \neq i$. This situation arises when a binary characteristic, such as employed/unemployed, is measured on subject i at time t . [Newton and Zhang \(1999, Sec. 6\)](#) consider an application where $X_{i,t} = 1$ if dairy cow i ’s milk sample at time t tests positive for pathogens and $X_{i,t} = 0$ otherwise. The relevant question here is whether the time series are homogeneous or, in other words, is there temporal dependence within the time series or are they independent? Below we construct a PRE-process and anytime valid test for independence versus a type of Markovian dependence.

Under the null hypothesis of time homogeneity, each X_i consists of a sequence of Bernoulli trials with its own success probability ω_i . In this case, the likelihood is binomial, which, for data $X^n = (X_1, \dots, X_n)$, is maximized at $\hat{\omega}_i = T^{-1} \sum_{t=1}^T X_{i,t}$, for $i = 1, \dots, n$. This maximum likelihood goes in the denominator of the PRE-process. For the alternative, we consider a mixture of Markov models where the mass function $q^\Psi(x)$ is given by [\(2\)](#) with kernel

$$p_u(x) = u_{00}^{t_{00}(x)} (1 - u_{00})^{t_{01}(x)} (1 - u_{11})^{t_{10}(x)} u_{11}^{t_{11}(x)}, \quad x = (x_1, \dots, x_T),$$

$t_{ab}(x)$ the number of transitions in series x from state $a \in \{0, 1\}$ to state $b \in \{0, 1\}$, and $u = (u_{00}, u_{11}) \in \mathbb{U} = [0, 1]^2$ is the pair of probabilities corresponding to the “transitions” $0 \rightarrow 0$ and $1 \rightarrow 1$, respectively. As described in [Section 3](#), we apply the PR algorithm to fit this mixture model and obtain likelihood $\hat{q}^{\text{PR}}(X^n)$, which is goes in the PRE-process’s numerator.

For illustration, we consider two simulation settings. The first is where the null is true, with each X_i being a sequence of Bernoulli trials with success probability $\omega_i \sim \text{Unif}(0, 1)$. The second is where the alternative is true, so that X_i is a Markov chain and its transition probability matrix is determined by the pair $(U_{00,i}, U_{11,i})$ iid $\text{Beta}(2, 4)$. In both cases, the processes are of length $T = 15$, and sample sizes are $n = 20, 40, \dots, 200$. The paths of $\log E_n^{\text{PR}}$ versus n for the null and alternative cases are shown in [Figure 3](#) and, as desired, we find that E_n^{PR} vanishes rapidly under the null and diverges similarly rapidly towards infinity under the alternative.




Figure 3: Plots of the PRE-process trajectories, $\log E_n^{\text{PR}}$ versus n , for the time homogeneity tests in Section 5.2.

6 Conclusion

This paper proposes a universal inference-inspired e-process construction based on the PR algorithm—a fast, recursive procedure for fitting flexible, nonparametric mixture models. We show that PRE-process is a genuine e-process and, therefore, offers provable, finite-sample, anytime valid inference. Further, we show that the PRE-process attains the optimal first-order growth rate under the alternative relative to the mixture model fit by PR. Numerical results demonstrate that the PRE-process’s finite-sample empirical performance closely agrees with the optimal behavior predicted by the asymptotic theory. Moreover, here and in the supplementary material, our proposal with a PR fit in the e-process numerator is shown to be computationally more efficient, numerically more stable, and more accurately approximates the oracle e-process compared to, say, a kernel density fit in the numerator.

The PRE-process developments here shed new light on PR and motivates further computational and theoretical investigations. In particular, there are open questions concerning PR’s convergence rates (as it pertains to our PRE-process’s growth rate here and otherwise) and a need for finite-sample bounds and efficient numerical methods for handling higher-dimensional multivariate data. On the last point, our PRticle filtering procedures employed here for multivariate problems is able to accommodate mixtures over moderate-dimensional space \mathbb{U} but, when the data and hence \mathbb{U} are high-dimensional, we do not expect the PRticle filter to be especially accurate, at least not in a reasonable time, so new ideas are needed to tackle those cases.

Acknowledgments

Thanks to Aaditya Ramdas for feedback on and corrections to a previous version, and to the other anonymous reviewers for their very helpful suggestions. This work is partially supported by the U.S. National Science Foundation, SES–2051225.

References

- Bernardo, J.-M. and Smith, A. F. M. (1994). *Bayesian Theory*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Ltd., Chichester.
- Blackwell, D. and Dubins, L. (1962). Merging of opinions with increasing information. *Ann. Math. Statist.*, 33:882–886.
- Bousquet, O. and Warmuth, M. K. (2002). Tracking a small set of experts by mixing past posteriors. *J. Mach. Learn. Res.*, 3:363–396.
- Choquet, G. (1953–1954). Theory of capacities. *Ann. Inst. Fourier, Grenoble*, 5:131–295 (1955).
- Cover, T. M. and Thomas, J. A. (2006). *Elements of information theory*. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition.
- Cule, M., Samworth, R., and Stewart, M. (2010). Maximum likelihood estimation of a multi-dimensional log-concave density. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(5):545–607.
- Darling, D. A. and Robbins, H. (1968). Some nonparametric sequential tests with power one. *Proc. Nat. Acad. Sci. U.S.A.*, 61:804–809.
- DasGupta, A. (2008). *Asymptotic Theory of Statistics and Probability*. Springer Texts in Statistics. Springer, New York.
- Dawid, A. P. (1984). Statistical theory. The prequential approach. *J. Roy. Statist. Soc. Ser. A*, 147(2):278–292.
- Dawid, A. P. (1985). Calibration-based empirical probability. *Ann. Statist.*, 13(4):1251–1285. With discussion.
- Dawid, A. P. (1991). Fisherian inference in likelihood and prequential frames of reference. *J. Roy. Statist. Soc. Ser. B*, 53(1):79–109. With discussion and a reply by the author.
- Denceux, T. (2019). Decision-making with belief functions: a review. *Internat. J. Approx. Reason.*, 109:87–110.
- Devaine, M., Gaillard, P., Goude, Y., and Stoltz, G. (2013). Forecasting electricity consumption by aggregating specialized experts: a review of the sequential aggregation of specialized experts, with an application to Slovakian and French country-wide one-day-ahead (half-)hourly predictions. *Mach. Learn.*, 90(2):231–260.
- Dixit, V. and Martin, R. (2023a). A PRticle filter algorithm for nonparametric estimation of multivariate mixing distributions. *Stat. Comput.*, 33(4):Paper No. 76, 14.
- Dixit, V. and Martin, R. (2023b). Revisiting consistency of a recursive estimator of mixing distributions. *Electron. J. Stat.*, 17(1):1007–1042.

- Doss, C. R. and Wellner, J. A. (2016). Global rates of convergence of the MLEs of log-concave and s -concave densities. *Ann. Statist.*, 44(3):954–981.
- Dubois, D. and Prade, H. (1988). *Possibility Theory*. Plenum Press, New York.
- Dümbgen, L. and Rufibach, K. (2011). logcondens: Computations related to univariate log-concave density estimation. *J. Stat. Softw.*, 39:1–28.
- Dunn, R., Gangrade, A., Wasserman, L., and Ramdas, A. (2024). Universal inference meets random projections: a scalable test for log-concavity. *J. Comput. Graph. Statist.*, to appear [arXiv:2111.09254](https://arxiv.org/abs/2111.09254).
- Duong, T. (2024). *ks: Kernel Smoothing*. R package version 1.14.2.
- Edwards, A. W. F. (1992). *Likelihood*. Johns Hopkins University Press, Baltimore, MD, expanded edition. Revised reprint of the 1972 original.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.*, 1:209–230.
- Gangrade, A., Rinaldo, A., and Ramdas, A. (2023). A sequential test for log-concavity. [arXiv:2301.03542](https://arxiv.org/abs/2301.03542).
- Ghosal, S. (2010). The Dirichlet process, related priors and posterior asymptotics. In *Bayesian Nonparametrics*, Camb. Ser. Stat. Probab. Math., pages 35–79. Cambridge Univ. Press, Cambridge.
- Ghosh, J. K. and Tokdar, S. T. (2006). Convergence and consistency of Newton’s algorithm for estimating mixing distribution. In Fan, J. and Koul, H., editors, *Frontiers in Statistics*, pages 429–443. Imp. Coll. Press, London.
- Grenander, U. (1956). On the theory of mortality measurement. II. *Skand. Aktuarietidskr.*, 39:125–153 (1957).
- Grünwald, P. (2022). Beyond Neyman–Pearson. [arXiv:2205.00901](https://arxiv.org/abs/2205.00901).
- Grünwald, P. (2023). The e-posterior. *Philos. Trans. Roy. Soc. A*, 381(2247):Paper No. 20220146, 21.
- Grünwald, P., de Heide, R., and Koolen, W. (2024). Safe testing. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, to appear; [arXiv:1906.07801](https://arxiv.org/abs/1906.07801).
- Hacking, I. (1976). *Logic of Statistical Inference*. Cambridge University Press, Cambridge–New York–Melbourne.
- Hahn, P. R., Martin, R., and Walker, S. G. (2018). On recursive Bayesian predictive distributions. *J. Amer. Statist. Assoc.*, 113(523):1085–1093.
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. (2021). Time-uniform, non-parametric, nonasymptotic confidence sequences. *Ann. Statist.*, 49(2):1055–1080.

- Jeffreys, H. (1998). *Theory of Probability*. Oxford Classic Texts in the Physical Sciences. The Clarendon Press, Oxford University Press, New York. Reprint of the 1983 edition.
- Kleijn, B. J. K. and van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *Ann. Statist.*, 34(2):837–877.
- Koenker, R. and Gu, J. (2017). Rebayes: An R package for empirical Bayes mixture methods. *J. Stat. Softw.*, 82(8):1–26.
- Kosorok, M. R. (2008). *Introduction to Empirical Processes and Semiparametric Inference*. Springer Series in Statistics. Springer, New York.
- Laird, N. (1978). Nonparametric maximum likelihood estimation of a mixed distribution. *J. Amer. Statist. Assoc.*, 73(364):805–811.
- Li, J. Q. and Barron, A. R. (2000). Mixture density estimation. In Solla, S., Leen, T., and Mueller, K.-R., editors, *Advances in Neural Information Processing Systems*, pages 279–285. MIT Press, Cambridge, Massachusetts.
- Liese, F. and Vajda, I. (1987). *Convex Statistical Distances*. Teubner, Leipzig.
- Lindsay, B. G. (1995). *Mixture Models: Theory, Geometry and Applications*. IMS, Hayward, CA.
- Littlestone, N. and Warmuth, M. K. (1994). The weighted majority algorithm. *Inform. and Comput.*, 108(2):212–261.
- Lo, A. Y. (1984). On a class of Bayesian nonparametric estimates. I. Density estimates. *Ann. Statist.*, 12(1):351–357.
- Mardia, K. V. and Jupp, P. E. (2000). *Directional Statistics*. Wiley Series in Probability and Statistics. John Wiley & Sons, Ltd., Chichester.
- Martin, R. (2021a). Inferential models and the decision-theoretic implications of the validity property. [arXiv:2112.13247](https://arxiv.org/abs/2112.13247).
- Martin, R. (2021b). A survey of nonparametric mixing density estimation via the predictive recursion algorithm. *Sankhya B*, 83(1):97–121.
- Martin, R. (2022). Valid and efficient imprecise-probabilistic inference with partial priors, II. General framework. [arXiv:2211.14567](https://arxiv.org/abs/2211.14567).
- Martin, R. (2023a). Comment on Waudby-Smith & Ramdas’s “Estimating means of bounded random variables by betting”. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, to appear; https://www4.stat.ncsu.edu/~rmartin/comment_rm.pdf.
- Martin, R. (2023b). A possibility-theoretic solution to Basu’s Bayesian–frequentist via media. *Sankhya A*, to appear, [arXiv:2303.17425](https://arxiv.org/abs/2303.17425).
- Martin, R. and Ghosh, J. K. (2008). Stochastic approximation and Newton’s estimate of a mixing distribution. *Statist. Sci.*, 23(3):365–382.

- Martin, R. and Liu, C. (2013). Inferential models: a framework for prior-free posterior probabilistic inference. *J. Amer. Statist. Assoc.*, 108(501):301–313.
- Martin, R. and Liu, C. (2015). *Inferential Models*, volume 147 of *Monographs on Statistics and Applied Probability*. CRC Press, Boca Raton, FL.
- Martin, R. and Tokdar, S. T. (2009). Asymptotic properties of predictive recursion: robustness and rate of convergence. *Electron. J. Stat.*, 3:1455–1472.
- Martin, R. and Tokdar, S. T. (2011). Semiparametric inference in mixture models with predictive recursion marginal likelihood. *Biometrika*, 98(3):567–582.
- Martin, R. and Tokdar, S. T. (2012). A nonparametric empirical Bayes framework for large-scale multiple testing. *Biostatistics*, 13(3):427–439.
- McVinish, R., Rousseau, J., and Mengersen, K. (2009). Bayesian goodness of fit testing with mixtures of triangular distributions. *Scand. J. Stat.*, 36(2):337–354.
- Newton, M., Kendzioriski, C., Richmond, C., Blattner, F., and Tsui, K. (2001). On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *J. Comput. Biology*, 8:37–52.
- Newton, M. A. (2002). On a nonparametric recursive estimator of the mixing distribution. *Sankhyā Ser. A*, 64(2):306–322.
- Newton, M. A., Quintana, F. A., and Zhang, Y. (1998). Nonparametric Bayes methods using predictive updating. In Dey, D., Müller, P., and Sinha, D., editors, *Practical Nonparametric and Semiparametric Bayesian Statistics*, volume 133 of *Lecture Notes in Statist.*, pages 45–61. Springer, New York.
- Newton, M. A. and Zhang, Y. (1999). A recursive algorithm for nonparametric analysis with missing data. *Biometrika*, 86(1):15–26.
- Neyman, J. and Pearson, E. S. (1933). On the problem of most efficient tests of statistical hypotheses. *Phil. Trans. Roy. Soc. London Ser. A*, 231:289–337.
- Neyman, J. and Scott, E. L. (1948). Consistent estimates based on partially consistent observations. *Econometrica*, 16:1–32.
- Patilea, V. (2001). Convex models, MLE and misspecification. *Ann. Statist.*, 29(1):94–123.
- Ramdas, A., Grünwald, P., Vovk, V., and Shafer, G. (2023). Game-theoretic statistics and safe anytime-valid inference. *Statist. Sci.*, 38(4):576–601.
- Ramdas, A., Ruf, J., Larsson, M., and W., K. (2022). Admissible anytime-valid sequential inference must rely on nonnegative martingales. [arXiv:2009.03167](https://arxiv.org/abs/2009.03167).
- Rissanen, J. (1984). Universal coding, information, prediction, and estimation. *IEEE Trans. Inform. Theory*, 30(4):629–636.

- Robbins, H. and Siegmund, D. (1971). A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics (Proc. Sympos., Ohio State Univ., Columbus, Ohio, 1971)*, pages 233–257. Academic Press, New York.
- Royall, R. M. (1997). *Statistical Evidence*, volume 71 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.
- Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, N.J.
- Shafer, G. (2021). Testing by betting: a strategy for statistical and scientific communication. *J. Roy. Statist. Soc. Ser. A*, 184(2):407–478.
- Shafer, G., Shen, A., Vereshchagin, N., and Vovk, V. (2011). Test martingales, Bayes factors and p -values. *Statist. Sci.*, 26(1):84–101.
- Shafer, G. and Vovk, V. (2019). *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, NJ.
- Tansey, W., Oluwasanmi, K., Poldrack, R. A., and Scott, J. G. (2018). False discovery rate smoothing. *J. Amer. Statist. Assoc.*, 113(523):1156–1171.
- Tao, H., Palta, M., Yandell, B. S., and Newton, M. A. (1999). An estimation method for the semiparametric mixed effects model. *Biometrics*, 55(1):102–110.
- Teicher, H. (1998). Strong laws for martingale differences and independent random variables. *J. Theoret. Probab.*, 11(4):979–995.
- Tokdar, S. T., Chakrabarti, A., and Ghosh, J. K. (2010). Bayesian nonparametric goodness of fit tests. In Cheh, M.-H., Dey, D., Müller, P., Sun, D., and Ye, K., editors, *Frontiers of Statistical Decision Making and Bayesian Analysis*, pages 185–194. Springer.
- Tokdar, S. T., Martin, R., and Ghosh, J. K. (2009). Consistency of a recursive estimate of mixing distributions. *Ann. Statist.*, 37(5A):2502–2522.
- Troffaes, M. C. M. and de Cooman, G. (2014). *Lower Previsions*. Wiley Series in Probability and Statistics. John Wiley & Sons, Ltd., Chichester.
- van de Geer, S. A. (2000). *Empirical Processes in M-Estimation*, volume 6 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes*. Springer-Verlag, New York.
- Vovk, V. G. (1993). A logic of probability, with application to the foundations of statistics. *J. Roy. Statist. Soc. Ser. B*, 55(2):317–351. With discussion and a reply by the author.

- Wald, A. (1945). Sequential tests of statistical hypotheses. *Ann. Math. Statistics*, 16:117–186.
- Wald, A. (1947). *Sequential Analysis*. John Wiley & Sons, Inc., New York.
- Wasserman, L., Ramdas, A., and Balakrishnan, S. (2020). Universal inference. *Proc. Natl. Acad. Sci. U.S.A.*, 117(29):16880–16890.
- Waudby-Smith, I. and Ramdas, A. (2024). Estimating means of bounded random variables by betting. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 86(1):1–27.
- Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Statist*, 9:60–62.
- Williamson, R. E. (1956). Multiply monotone functions and their Laplace transforms. *Duke Math. J.*, 23:189–207.
- Woody, S., Padilla, O. H. M., and Scott, J. G. (2022). Optimal post-selection inference for sparse signals: a nonparametric empirical Bayes approach. *Biometrika*, 109(1):1–16.

Supplementary Material

This supplementary material contains some additional technical details that could not fit in the main paper. In particular, Section 1 offers some new perspectives on the predictive recursion (PR) algorithm in terms of Dawid’s prequential statistics; Section 2 offers some further details on the new PRE-process proposed in the main paper; Section 3 gives a detailed statement of the conditions needed to analyze the PR estimates, along with a proof of the asymptotic growth rate result (Theorem 2) in the main paper; Section 4 present some old and new details on the rate of convergence of PR’s estimator; Section 5 offers some empirical results to help pin down the magnitude of the “ $o(n)$ ” error term in Theorem 2’s asymptotic approximation; Section 6 provides a comparison of two e-processes in multivariate settings; and Section 7 proposes a variation on the PRE-process presented in the main paper.

A New perspectives on PR

Both of us wrote our PhD theses on predictive recursion (PR), viewed as an estimation procedure. That is, the focus of those previous efforts was on large-sample consistency of PR’s estimated mixing and mixture distributions and on applications in high-dimensional inference. The present paper takes a different perspective, focusing on PR in the context of *prediction*, which, ironically, has not been thoroughly investigated in the literature. This is both interesting and relevant to the main paper, so we collect some of these developing thoughts—potential new lines of research—here in this supplementary material. We specifically would like to thank one of the anonymous reviewers for pointing out some of these connections.

First, there are a number of methods available in the machine learning literature for online or sequential prediction, i.e., algorithms that produce a sequence of probability distributions (Q_t) , $t = 1, 2, \dots$, where Q_t is interpreted as the predictive distribution of X_{t+1} depending on X^t . Like with PR, it is not uncommon to construct these predictive distributions using mixtures. For example, in [Devaine et al. \(2013\)](#), the authors consider applications where there is a fixed, finite number of experts making forecasts about X_{t+1} based on previous data X^t , and the proposed algorithm aggregates the experts’ forecasts using mixtures; see also [Littlestone and Warmuth \(1994\)](#) and [Bousquet and Warmuth \(2002\)](#). [Devaine et al. \(2013\)](#) go on to prove non-asymptotic bounds on the regret—prediction performance compared to an oracle—for certain convex loss functions. While the particular algorithms in the above references are different from PR, this connection is relevant for at least two reasons: first, any sequential aggregation algorithm ought to have its own corresponding e-process—an AGGREGEE-process, say?—that warrants investigation, just like PR and our proposed PRE-process; second, perhaps the theoretical analysis of these alternative algorithms can shed light on, say, how to establish similar regret bounds for PR.

Next, and particularly relevant to the details in the present paper, is a connection to Dawid’s *prequential* approach to and perspective on probability and statistics (e.g., [Dawid 1984, 1991](#)). Even the names are similar: the “prequential” is a combination of the words “predictive” and “sequential,” which clearly has aspects in common with

“predictive recursion.” A key difference, however, is that what Dawid has put forward is a general theory or perspective, while PR provides a particular instantiation of it, as we now explain. The sequence (Q_t) discussed in the previous paragraph is what Dawid calls a *prequential forecasting system*, which could take all sorts of different forms, e.g., from a fully specified joint distribution for the observables to a more ad hoc expert aggregation strategy. PR falls somewhere in between these two extremes: PR has at its core a flexible, nonparametric mixture model for the observables, but it learns from observations and updates in a non-Bayesian/non-probabilistic fashion. Dawid argues, as formulated in his *prequential principle*, that the quality of a forecasting system should be determined solely based on how well it predicts the actual observations; in other words, this assessment should not take into consideration hypothetical outcomes that were not observed. Our proposed PRE-process test in the main paper is exactly in this spirit. Indeed, just as in Dawid (1984, Sec. 7), our PRE-process in Equation (6) in the main text, is the ratio of two forecasting systems, each applied to the observed data sequence: the numerator represents the PR-based forecasts and the denominator a summary of the forecasts under the null model \mathcal{P}_0 . Then the proposed test rejects the null model if and only if the ratio exceeds some specified threshold (which does not depend on, say, “null distribution” characteristics, etc). In other words, the null model is rejected if and only if the PR-based forecasting system is significantly “better” than that which assumes the null model. That our PRE-process test is consistent with the prequential principle itself is not enough to justify its use. It is precisely for this reason that Theorems 1–2 in the main paper are essential: on the one hand, the PRE-process (E_n^{PR}) stays small forever when \mathcal{P}_0 is true and, on the other hand, (under certain conditions) it diverges quickly to infinity when \mathcal{P}_0 is not true.

Dawid (1985) discusses the use of forecasting systems to define an objective, empirical notion of probability assigned to observed data sequences, and comments on the implications to the foundations of statistics, probability, and science more broadly; see, also, Vovk (1993). Not just any forecasting system achieves this objective, and Dawid highlights a key *calibration* condition. Using our present notation, fix a sequence (A_n) of subsets of \mathbb{X} and the observations $a_n = 1(X_n \in A_n)$, where $1(\cdot)$ is the indicator. Sparing some technical details—concerning computable selection rules, essential to properly defining “randomness” in the spirit of von Mises—a forecasting system (Q_n) is *calibrated* with respect to the observed sequence (a_n) if

$$\frac{1}{n} \sum_{t=1}^n \{Q_t(A_t) - a_t\} \rightarrow 0, \quad n \rightarrow \infty.$$

This has important implications, namely, that any two forecasting systems that are calibrated to (a_n) must be asymptotically equivalent (e.g., Blackwell and Dubins 1962), and it is precisely that common “asymptotic probability” that Dawid puts forward as the empirical probability of (a_n) . For the PR forecasting system (\hat{Q}_n^{PR}) , if X_1, X_2, \dots are iid

P^* , then

$$\begin{aligned} \frac{1}{n} \sum_{t=1}^n \{\widehat{Q}_t^{\text{PR}}(A_t) - a_t\} &= \frac{1}{n} \sum_{t=1}^n \{\widehat{Q}_t^{\text{PR}}(A_t) - P^*(A_t) + P^*(A_t) - a_t\} \\ &= \frac{1}{n} \sum_{t=1}^n \{\widehat{Q}_t^{\text{PR}}(A_t) - P^*(A_t)\} + \frac{1}{n} \sum_{t=1}^n \{P^*(A_t) - a_t\}. \end{aligned}$$

It follows from the standard law of large numbers for bounded, mean zero random variables that the second term in the last line above converges to 0 with P^* -probability 1. If P^* is contained in the set \mathcal{Q} defined in Section 3.3 of the main paper, and if the conditions in Section 3 below are satisfied, then it follows from Corollary 4.6 in [Martin and Tokdar \(2009\)](#) that $\widehat{Q}_n^{\text{PR}} \rightarrow P^*$ with P^* -probability 1 in total variation. This implies that the first term in the above display is vanishing as $n \rightarrow \infty$ and, consequently, the PR forecasting system is calibrated with respect to the sequence $a_n = 1(X_n \in A_n)$ determined by X^∞ , for any such X^∞ in a set of P^* -probability 1. This, of course, is an especially simple case; but PR itself is quite flexible, e.g., it does not make use of the iid assumption, so we suspect that calibration can be established more generally.

Finally, under this umbrella of forecasting systems, the idea put forward in [Martin and Tokdar \(2011\)](#) can be given a new, potentially useful twist. They proposed a semi-parametric mixture model wherein the kernel $\{p_u : u \in \mathbb{U}\}$ used to define the PR algorithm depends on some parameter $\theta \in \mathbb{T}$; that is, their kernel is written as $p_{u,\theta}$, with $x \mapsto p_{u,\theta}(x)$ a probability density function for each (u, θ) . At any fixed value of θ , we can produce PR's joint marginal density for X^n , denoted here as $\hat{q}_\theta^{\text{PR}}(X^n)$. Then [Martin and Tokdar \(2011\)](#) defined the ‘‘likelihood’’ function $\theta \mapsto \hat{q}_\theta^{\text{PR}}(X^n)$ and compared it to the marginal likelihood function under a Bayesian nonparametric mixture model with a Dirichlet process prior. They then proposed an estimate of θ based on maximizing that PR-marginal likelihood: $\hat{\theta}_n^{\text{PR}} = \arg \max_\theta \hat{q}_\theta^{\text{PR}}(X^n)$. Unfortunately, despite ample empirical evidence suggesting it, sufficiently general results on large-sample consistency of $\hat{\theta}_n^{\text{PR}}$ have yet to be established. But one can interpret Martin & Tokdar's suggestion as a particular instantiation of the ‘‘prequential model’’ in [Dawid \(1991, Sec. 7.3\)](#) which, for each θ , evaluates forecasts for the next observation in light of what has already been observed. If, as conjectured in the previous paragraph, this PR-based prequential model can be shown to have high-level justification in terms of calibration, etc., then it would be meaningful to assess the properties of $\hat{\theta}_n^{\text{PR}}$ relative to the PR-based model. This creates a new and different opportunity to demonstrate the utility of this estimator. Indeed, [Dawid \(1984, 1991\)](#) presents a notion of (simple) prequential consistency, in which the predictions of X_{n+1} based on $\widehat{Q}_{\hat{\theta}_n^{\text{PR}}}^{\text{PR}}$ and $\widehat{Q}_{\theta^*}^{\text{PR}}$, given $X^n = x^n$, are compared as $n \rightarrow \infty$, where θ^* denotes the ‘‘true’’ value in the PR-based prequential model. If these predictions are asymptotically equivalent, then Dawid would call $\hat{\theta}_n^{\text{PR}}$ (simply and prequentially) *consistent*. Aside from the attractiveness of the prequential perspective, we believe that the aforementioned demonstration is within reach. This is also promising too, as Dawid argues, this notion of consistency often holds more broadly than the classical notions of consistency.

B Further PRe-process details

The paper’s main text describes and focuses on the use of our proposed PRe-process E_n^{PR} for anytime valid hypothesis testing. But there is more that can be done with the PRe-process and here we briefly describe three such developments.

First, given a null model, \mathcal{P}_0 , one can define a corresponding “anytime p-value”

$$\pi^{\text{PR}}(X^n; \mathcal{P}_0) = 1 \wedge E^{\text{PR}}(X^n; \mathcal{P}_0)^{-1}.$$

Just like in Corollary 1 in the main paper, it follows that

$$\sup_{P \in \mathcal{P}_0} P\{\pi^{\text{PR}}(X^N; \mathcal{P}_0) \leq \alpha\} \leq \alpha, \quad \text{all } \alpha \in [0, 1] \text{ and all stopping times } N.$$

So, as expected, the anytime valid test proposed in the main paper is equivalent to a test that rejects when the above p-value is no more than α .

Second, by considering singleton null hypotheses $\mathcal{P}_0 = \{P_0\}$ and varying P_0 , the above p-value defines a function $P_0 \mapsto \pi^{\text{PR}}(X^n; \{P_0\})$ which, as explained in [Martin (2023a)], determines a data-dependent imprecise probability distribution supported on \mathcal{P} , in particular, a necessity–possibility measure pair (e.g., [Dubois and Prade 1988; Shafer 1976]). This, in turn, determines an anytime valid *inferential model* ([Martin 2022; Martin and Liu 2013, 2015]) that provides provably safe and reliable imprecise-probabilistic uncertainty quantification about the unknown P . Akin to a Bayesian posterior distribution, this imprecise-probabilistic uncertainty quantification can then be used, e.g., for formal decision-making ([Martin 2021a]). Indeed, if $(a, P) \mapsto \ell_a(P)$ quantifies the loss of taking action a when the state of the world is P , then one can find a lower and upper expected loss using a specialization of Choquet’s general theory of integration with respect to capacities ([Choquet 1954]) to the case of necessity and possibility measures (e.g., [Troffaes and de Cooman 2014, Sec. 7.8]). From here, one can proceed to, say, minimize the upper expected loss with respect to the action a to obtain an inferential model version of a Bayes rule, and establish bounds on the difference between the inferential model’s upper expected loss and the oracle loss that knows the true P . Since this is not specific to the PRe-process developments in this paper, we will not go into any further details here. But the interested reader might like to see [Dencœux (2019)] for a review of the literature on this brand of generalized Bayes decision theory, and also R. Martin’s extended abstract in the forthcoming Oberwolfach Mathematical Research Institute report (<https://www.mfo.de/occasion/2419b>).

Finally, a more familiar consequence of the anytime validity of the PRe-process-based test in the main paper is inverting the test to construct an anytime valid confidence set. Let $\phi : \mathcal{P} \rightarrow \Phi$ be a map that extracts some relevant feature from $P \in \mathcal{P}$. For example, $\phi(P) = \int_{\mathbb{X}} x P(dx)$ is the mean of P , $\phi(P) = \inf\{x : P((-\infty, x]) \geq \tau\}$ is the τ^{th} quantile of P , and $\phi(P) = P$ is the distribution P itself. For a given ϕ , define the following data-dependent subset of Φ :

$$C_\alpha(X^n) = \{\phi(P) : \pi^{\text{PR}}(X^n; \{P\}) > \alpha\}, \quad \alpha \in [0, 1].$$

The set $C_\alpha(X^n)$ is just the collection of all $\phi(P)$ corresponding to “null” P s that the PRe-process-based test would not reject for data X^n at level α . Then the fact that E_n^{PR} is

an e-process implies that $C_\alpha(\cdot)$ is an anytime valid $100(1 - \alpha)\%$ confidence set for $\phi(P)$ in the sense that

$$\sup_{P \in \mathcal{P}} P\{C_\alpha(X^N) \not\supseteq \phi(P)\} \leq \alpha, \quad \text{any } \alpha \in (0, 1) \text{ and any stopping time } N.$$

Of course, identifying the set $C_\alpha(X^n)$ for a given ϕ and a given data set X^n could be a challenge, but approximations could surely be found if there was sufficient desire to do so.

C The PRe-process's asymptotic growth rate

C.1 Setup of Theorem 2

There are a few inputs that play key roles in the asymptotic properties of the PR algorithm. These include some user-specified inputs, namely, the family of kernel densities $\{p_u : u \in \mathbb{U}\}$ that determine the set of mixtures \mathcal{Q} , and PR's weights (w_i) and initial guess Ψ_0 ; two more relevant quantities that are beyond the user's control are the true distribution P^* , with density p^* , and the corresponding Kullback–Leibler projection of P^* onto the set \mathcal{Q} of mixtures. This latter projection is defined as the distribution Q^* with density q^* that satisfies

$$K(P^*, Q^*) = K(P^*, \mathcal{Q}) := \inf_{Q \in \mathcal{Q}} K(P^*, Q). \quad (12)$$

The existence of Q^* is ensured by various sets of conditions; see, e.g., [Liese and Vajda \(1987, Ch. 8\)](#). In particular, existence of Q^* is implied by Condition [1](#) below, as shown in Lemma 3.1 of [Martin and Tokdar \(2009\)](#). The following conditions, on which Theorem 2 is based, concern the basic properties of these individual inputs and their interplay.

Condition 1. \mathbb{U} is compact and $u \mapsto p_u(x)$ is continuous for almost all x .

Condition 2. The PR weight sequence (w_i) satisfies

$$\sum_{i=1}^{\infty} w_i = \infty \quad \text{and} \quad \sum_{i=1}^{\infty} w_i^2 < \infty. \quad (13)$$

Condition 3. The kernel density $p_u(x)$ satisfies

$$\sup_{u_1, u_2 \in \mathbb{U}} \int \left\{ \frac{p_{u_1}(x)}{p_{u_2}(x)} \right\}^2 p^*(x) dx < \infty \quad (14)$$

Condition 4. The Kullback–Leibler projection Q^* in [\(12\)](#), with density q^* , satisfies

$$\int \left\{ \log \frac{p^*(x)}{q^*(x)} \right\}^2 p^*(x) dx < \infty. \quad (15)$$

Here we offer some explanation and intuition. First, compactness of the mixing distribution support \mathbb{U} in Condition [1](#) is difficult to relax, but the fact that \mathbb{U} can be taken arbitrarily large means that this imposes effectively no practical constraints on the user.

Continuity of the kernel can be relaxed, but at the expense of a much more complicated condition; the reader interested in this can consult Equation (7) in [Dixit and Martin \(2023b\)](#) and the relevant discussion. Condition [2](#) says simply that the weights must be vanishing to ensure convergence but not too quickly since the algorithm needs an opportunity to learn; the requirement in [\(13\)](#) is just right for this. Condition [3](#) is non-trivial but holds for Gaussian and other exponential family distributions thanks to the compactness of \mathbb{U} in Condition [1](#). Finally, Condition [4](#) concerns the quality of the mixture model itself. One cannot hope to achieve quality estimation/inference in any sense if the “best” member of the mixture model differs considerably from the true density p^* . Equation [\(15\)](#) is just a particular way to say that p^* and q^* do not differ by too much.

Once the user makes his/her specification of the mixture model, Condition [3](#) determines a set of true densities p^* for which the PR algorithm will provide consistent estimation. Indeed, Theorem 1 in [Martin and Tokdar \(2011\)](#) states that, under Conditions [1-3](#), the PR estimator \hat{q}_{X^n} satisfies $K(p^*, \hat{q}_{X^n}) \rightarrow \inf_{q \in \mathcal{Q}} K(p^*, q)$ with P^* -probability 1 as $n \rightarrow \infty$. The user, of course, can vary the mixture model specification to tailor PR toward what they expect p^* to look like. But we need more than consistency for our purposes here, and Conditions [4](#) further restricts the set of true densities to those for which the PR algorithm can give us the “more” that we need.

C.2 Proof of Theorem 2

As explained in the main text, Theorem 2 is a consequence of the following property of PR:

$$n^{-1} \log \hat{q}^{\text{PR}}(X^n) = n^{-1} \log p^*(X^n) - K(P^*, \mathcal{Q}) + o(1), \quad n \rightarrow \infty, \quad (16)$$

The goal here is simply to show that [\(16\)](#) holds under Conditions [1-4](#) as stated above. The argument closely follows that in [Martin and Tokdar \(2011\)](#), but we provide the details here for completeness since their context and notation is different from ours.

The strategy of the proof is as follows. First, simplify the notation by writing $\hat{q}_{i-1}(X_i) = \hat{q}_{X^{i-1}}(X_i)$ for each i . Next, define the sequence of random variables

$$K_n = \frac{1}{n} \sum_{i=1}^n \log \frac{p^*(X_i)}{\hat{q}_{i-1}(X_i)},$$

which might be interpreted as a sort of empirical Kullback–Leibler divergence. If $K^* = K(P^*, \mathcal{Q})$, then [\(16\)](#) is equivalent to $K_n \rightarrow K^*$ with P^* -probability 1 as $n \rightarrow \infty$. It is this latter claim that we will prove below, using a martingale strong law in [Teicher \(1998\)](#).

Towards this, define a sequence of random variables Z_i as,

$$Z_i = \log \frac{p^*(X_i)}{\hat{q}_{i-1}(X_i)} - K(p^*, \hat{q}_{i-1}), \quad i \geq 1.$$

Recall that $\mathcal{A}_{i-1} = \sigma(X^{i-1})$, so $E_{P^*}(Z_i | \mathcal{A}_{i-1}) = 0$ and, therefore, $\{(Z_i, \mathcal{A}_i) : i \geq 1\}$ is a

zero mean martingale sequence under P^* . For q^* as in (12), we have

$$\begin{aligned}
E_{P^*}(Z_i^2 \mid \mathcal{A}_{i-1}) &\leq \int \left\{ \log \frac{p^*(x)}{\hat{q}_{i-1}(x)} \right\}^2 p^*(x) dx \\
&= \int \left\{ \log \frac{q^*(x)}{\hat{q}_{i-1}(x)} + \log \frac{p^*(x)}{q^*(x)} \right\}^2 p^*(x) dx \\
&\leq 2 \int \left\{ \log \frac{q^*(x)}{\hat{q}_{i-1}(x)} \right\}^2 p^*(x) dx + 2 \int \left\{ \log \frac{p^*(x)}{q^*(x)} \right\}^2 p^*(x) dx \\
&= 2A_i + 2B.
\end{aligned}$$

Of course, B is a finite constant according to Condition 4. To bound A_i let us first define $\mathbb{X}_0 = \{x : q^*(x) < \hat{q}_{i-1}(x)\}$. Using basic properties of the logarithm, we get

$$\begin{aligned}
A_i &= \int \left\{ \log \frac{q^*(x)}{\hat{q}_{i-1}(x)} \right\}^2 p^*(x) dx \\
&= \int_{\mathbb{X}_0} \left\{ \log \frac{\hat{q}_{i-1}(x)}{q^*(x)} \right\}^2 p^*(x) dx + \int_{\mathbb{X}_0^c} \left\{ \log \frac{q^*(x)}{\hat{q}_{i-1}(x)} \right\}^2 p^*(x) dx \\
&\leq \int_{\mathbb{X}_0} \left\{ \frac{\hat{q}_{i-1}(x)}{q^*(x)} - 1 \right\}^2 p^*(x) dx + \int_{\mathbb{X}_0^c} \left\{ \frac{q^*(x)}{\hat{q}_{i-1}(x)} - 1 \right\}^2 p^*(x) dx \\
&\leq 2 + \int \left[\left\{ \frac{\hat{q}_{i-1}(x)}{q^*(x)} \right\}^2 + \left\{ \frac{q^*(x)}{\hat{q}_{i-1}(x)} \right\}^2 \right] p^*(x) dx.
\end{aligned}$$

Since both \hat{q}_{i-1} and q^* in the two numerators in the above display are mixtures of the kernel p_u , we can say that

$$A_i \leq 2 + 2 \sup_{u_1, u_2 \in \mathbb{U}} \int \left\{ \frac{p_{u_1}(x)}{p_{u_2}(x)} \right\}^2 p^*(x) dx$$

which is bounded by Condition 3. Therefore, $E_{P^*}(Z_i^2 \mid \mathcal{A}_{i-1})$ is bounded too. Switching from index “ i ” to the more natural “ n ,” we have that

$$\frac{E_{P^*}(Z_n^2 \mid \mathcal{A}_{n-1})}{n^2(\log \log n)^{-1}} \lesssim n^{-2}(\log \log n) \rightarrow 0.$$

Therefore, by Markov’s inequality, with P^* -probability 1 we have,

$$\sum_{n=1}^{\infty} P^* \left(|Z_n| > \frac{n}{\log \log n} \mid \mathcal{A}_{n-1} \right) \lesssim \sum_{n=1}^{\infty} \frac{(\log \log n)^2}{n^2} < \infty.$$

From this and Corollary 2 of Teicher (1998)—with his “ $\beta = 1$ ”—we get $n^{-1} \sum_{i=1}^n Z_i \rightarrow 0$ with P^* -probability 1. Therefore, also with P^* -probability 1,

$$\left| K_n - \frac{1}{n} \sum_{i=1}^n K(p^*, \hat{q}_{i-1}) \right| = \left| (K_n - K^*) - \frac{1}{n} \sum_{i=1}^n \{K(p^*, \hat{q}_{i-1}) - K^*\} \right| \rightarrow 0.$$

From Theorem 1 in Martin and Tokdar (2011) and Cesaro’s theorem, we have $K_n - K^* \rightarrow 0$ with P^* -probability 1.

C.3 Condition (9) concerning the null model

Theorem 2 in the main paper—see Equation (9)—relies on a particular feature $\kappa^*(\mathcal{P}_0)$ of the null model \mathcal{P}_0 , namely,

$$\liminf_{n \rightarrow \infty} n^{-1} \log\{p^*(X^n)/\hat{p}_0(X^n)\} \geq \kappa^*(\mathcal{P}_0), \quad \text{with } P^*\text{-probability 1,} \quad (17)$$

where $\hat{p}_0(X^n) = \sup_{p_0 \in \mathcal{P}_0} p_0(X^n)$ is the maximum likelihood in \mathcal{P}_0 . As stated in the main paper, it is easy to check that $\kappa^*(\mathcal{P}_0)$ is bounded above by $K(P^*, \mathcal{P}_0)$. Indeed, by definition of maximum likelihood,

$$n^{-1} \log\{p^*(X^n)/\hat{p}_0(X^n)\} \leq n^{-1} \log\{p^*(X^n)/p_0(X^n)\}, \quad \text{for all } P_0 \in \mathcal{P}_0,$$

and the upper bound converges with P^* -probability 1 to $K(P^*, P_0)$. This can be made arbitrarily close to $K(P^*, \mathcal{P}_0)$ through careful choice of P_0 and, therefore, $\kappa^*(\mathcal{P}_0) \leq K(P^*, \mathcal{P}_0)$.

Our plan here is to focus on the case where $\kappa^*(\mathcal{P}_0) = K(P^*, \mathcal{P}_0)$, but we first want to emphasize that strict inequality is possible. This corresponds to a particularly unusual case where the maximum likelihood estimator fits the data better, even asymptotically, than the true distribution P^* (or the model’s “best representative”). As an illustration, consider the famous Neyman–Scott problem (Neyman and Scott 1948), where data X^n is of the form $\{(X_{11}, \dots, X_{1J}), \dots, (X_{n1}, \dots, X_{nJ})\}$, fully independent throughout, with X_{i1}, \dots, X_{in} iid $\mathcal{N}(\mu_i, \sigma^2)$, for $i = 1, \dots, n$. The point is that there are n -many mean parameters (μ_1, \dots, μ_n) but only one variance parameter σ^2 . The maximum likelihood estimators are

$$\hat{\mu}_i = \frac{1}{J} \sum_{j=1}^J X_{ij}, \quad i = 1, \dots, n \quad \text{and} \quad \hat{\sigma}^2 = \frac{1}{nJ} \sum_{i=1}^n \sum_{j=1}^J (X_{ij} - \hat{\mu}_i)^2.$$

What makes this example “famous” is that it is one of the first to demonstrate inconsistency of the maximum likelihood estimator. In particular, if σ^2 denotes the true variance parameter under P^* , then $\hat{\sigma}^2 \rightarrow \frac{J-1}{J}\sigma^2 < \sigma^2$ with P^* -probability 1 as $n \rightarrow \infty$ with J fixed. In this case, using the general notation from above, it can be shown that

$$n^{-1} \log\{p^*(X^n)/\hat{p}_0(X^n)\} \rightarrow \log\left(\frac{J-1}{J}\right) < 0,$$

while the minimum Kullback–Leibler number for this correctly-specified model is 0. This happens because maximum likelihood systematically underestimates the variance, and a model with less spread around the observations will tend to fit better in the sense of having non-trivially higher likelihood. Our hunch is that strict inequality $\kappa^*(\mathcal{P}_0) < K(P^*, \mathcal{P}_0)$ holds only when the \mathcal{P}_0 -maximum likelihood estimator is inconsistent, but we will not pursue this matter here.

Returning to the general question, we are interested in when

$$n^{-1} \log\{p^*(X^n)/\hat{p}_0(X^n)\} \rightarrow K(P^*, \mathcal{P}_0), \quad \text{with } P^*\text{-probability 1.}$$

Of course, if the model is correct in the sense that $P^* \in \mathcal{P}_0$, then the right-hand side is 0 and we would expect that the left-hand would be a limit that also equals 0. Corollary 3.2

in [Doss and Wellner \(2016\)](#) gives a result along these lines for the case of log-concave densities, which was one of the applications in the main paper. The more interesting case for us is when the model is wrong, i.e., $K(P^*, \mathcal{P}_0) > 0$. As is common in the literature on model misspecification (e.g., [Kleijn and van der Vaart 2006](#); [Patilea 2001](#)), we assume that there exists $P_0^\dagger \in \mathcal{P}_0$ such that $K(P^*, P_0^\dagger) = K(P^*, \mathcal{P}_0)$. Uniqueness of P_0^\dagger is not necessary for us here, but it is unique if \mathcal{P}_0 is convex, e.g., in the monotone density application in the main paper. Then the scaled log-likelihood ratio can be written as

$$n^{-1} \log\{p^*(X^n)/\hat{p}_0(X^n)\} = n^{-1} \log\{p^*(X^n)/p_0^\dagger(X^n)\} + n^{-1} \log\{p_0^\dagger(X^n)/\hat{p}_0(X^n)\}.$$

The law of large numbers implies that the first term on the right-hand side converges to $K(P^*, \mathcal{P}_0)$ with P^* -probability 1. So, only the second term on the right-hand side requires further analysis and, in particular, we would like to know when this would be vanishing as $n \rightarrow \infty$. Since $\hat{p}_0(X^n) \geq p_0^\dagger(X^n)$, it is enough to show that $n^{-1} \log\{\hat{p}_0(X^n)/p_0^\dagger(X^n)\} \rightarrow 0$.

A general strategy is through the use of a suitable uniform law of large numbers. Consider the following decomposition, with de Finetti's notation P^*f for expectation $\int f dP^*$,

$$\begin{aligned} n^{-1} \log\{\hat{p}_0(X^n)/p_0^\dagger(X^n)\} &= \sup_{P_0 \in \mathcal{P}_0} \frac{1}{n} \sum_{i=1}^n \{\log p_0(X_i) - \log p_0^\dagger(X_i)\} \\ &= \sup_{P_0 \in \mathcal{P}_0} \frac{1}{n} \left[\sum_{i=1}^n \{\log p_0(X_i) - P^* \log p_0\} \right. \\ &\quad \left. - \sum_{i=1}^n \{\log p_0^\dagger(X_i) - P^* \log p_0^\dagger\} \right. \\ &\quad \left. - n\{K(P^*, P_0) - K(P^*, P_0^\dagger)\} \right] \\ &\leq \sup_{P_0 \in \mathcal{P}_0} \left| \frac{1}{n} \sum_{i=1}^n \{\log p_0(X_i) - P^* \log p_0\} \right| + o(1), \end{aligned}$$

where “ $o(1)$ ” corresponds to $n^{-1} \sum_{i=1}^n \{\log p_0^\dagger(X_i) - P^* \log p_0^\dagger\}$, which is vanishing with P^* -probability 1 by the law of large numbers. Of course, if $\log \mathcal{P}_0$ is a Glivenko–Cantelli class, then of course the upper bound in the above display is vanishing and the desired result holds. If, for example, the densities in \mathcal{P}_0 are not bounded away from 0, then this naive Glivenko–Cantelli property might fail. In that case, one can restrict the class \mathcal{P}_0 to some better-behaved subset for which the Glivenko–Cantelli property can be established, then show that the maximum likelihood estimator would be in that restricted \mathcal{P}_0 with probability converging to 1. This is relatively standard from here, so we refer the interested reader to the classical texts on empirical process theory and statistical inference/learning (e.g., [Kosorok 2008](#); [van de Geer 2000](#); [van der Vaart 1998](#); [van der Vaart and Wellner 1996](#)).

D On PR's rate of convergence

The rate at which PR's density estimator converges to the true density p^* is of general interest and is relevant to questions about the asymptotic growth rate of the PRE-process

under the alternative. Here we review and slightly extend the theoretical analysis in [Martin and Tokdar \(2009\)](#).

[Martin and Tokdar \(2009\)](#) showed that, under the conditions of Theorem 2 in the main paper (see Section [C](#) above), if P^* is in the interior of the set \mathcal{Q} of mixtures and if PR’s weight sequence (w_i) has the form $w_i = (i + 1)^{-\gamma}$, for $\gamma \in (2/3, 1]$, then PR’s density estimator \hat{q}_{X^n} satisfied $n^{1-\gamma}K(p^*, \hat{q}_{X^n}) \rightarrow 0$ in P^* -probability. It is for precisely this reason that we take $\gamma = 0.67$, just slightly more than $2/3$, in our numerical examples—it corresponds to roughly the best rate according to the available theory. Martin & Tokdar argue that their bound on the rate is conservative or worst-case in the sense that it corresponds to a P^* arbitrarily close to the boundary of \mathcal{Q} ; this boundary case is “worst” because PR struggles most to estimate a mixing distribution that does not have a density with respect to the posited dominating measure.

As an aside, a shortcoming of the theoretical analyses of PR published to date is that they apparently are not able to accommodate structural assumptions about P^* , so the results all are of a “worst-case” form. This makes it difficult to assess the performance of PR in natural cases where P^* has a sufficiently smooth density, etc.

It turns out that there is another previously-unknown sense in which PR’s convergence rate is conservative. The following proposition makes this precise.

Proposition 1. *Under the conditions stated in the second paragraph of this section,*

$$n^{1-\gamma} \max_{s \geq n} K(p^*, \hat{q}_{X^s}) \rightarrow 0 \quad \text{in } P^*\text{-probability, } n \rightarrow \infty.$$

Proof. Martin & Tokdar’s proof leans on a demonstration that $K(p^*, \hat{q}_{X^n})$ and $n^{1-\gamma}K(p^*, \hat{q}_{X^n})$ are what [Robbins and Siegmund \(1971\)](#) call an “almost supermartingale” sequences, and then they apply Robbins & Siegmund’s Theorem 1 to prove that both are vanishing with P^* -probability 1. But Robbins & Siegmund also prove a maximal inequality for almost supermartingales and the above proposition’s claim is a consequence of their Proposition 2 and Martin & Tokdar’s Theorem 4.8. We omit the notation-heavy details here. \square

This result implies that Martin & Tokdar’s conservative $n^{-(1-\gamma)}$ rate for $K(p^*, \hat{q}_{X^n})$ is more accurately the rate associated with the obviously-no-smaller quantity $\max_{s \geq n} K(p^*, \hat{q}_{X^s})$. This is non-trivial because, like supermartingales, almost supermartingale sequences only “tend” to decrease. Unfortunately, it is not clear how this new sense of conservatism in the extant theoretical analyses of PR can be used to pin down a less conservative bound on rate of convergence. But the empirical work below may shed some light on this.

E Size of the error term in Theorem 2

A question posed in the main paper is: how large is the PR-process’s “ $o(n)$ ” term in Theorem 2? The theoretical analysis in [Martin and Tokdar \(2009\)](#) and Section [D](#) above imply a conservative bound of n^γ on the “ $o(n)$ ” term in Theorem 2 of the main paper, when PR’s weight sequence (w_i) satisfies $w_i = (i + 1)^{-\gamma}$, for $\gamma \in (2/3, 1]$. So, with the suggested choice of $\gamma \approx 2/3$, a conservative bound on the “ $o(n)$ ” term would be roughly $n^{2/3}$. Less conservatively, away from the edge cases, we expect an improved rate, and the current conjecture is that the $o(n)$ term in question is roughly $n^{\gamma/2}$, or roughly $n^{1/3}$.

when $\gamma \approx 2/3$. In this section, we present some simulation results as empirical evidence in support of this conjecture.

Towards this, recall the oracle e-process defined in the main paper as

$$E_n^{\text{OR}} = \frac{p^*(X^n)}{\sup_{P_0 \in \mathcal{P}_0} p_0(X^n)},$$

whose magnitude can be used to compare the null model \mathcal{P}_0 to the true distribution P^* based on the data X^n . Then the difference between the log-oracle e-process and the log-PR-process is

$$D_n := \log E_n^{\text{OR}} - \log E_n^{\text{PR}} = \sum_{i=1}^n \log \frac{p^*(X_i)}{\hat{q}_{X^{i-1}}(X_i)},$$

which is easy to evaluate numerically and does not depend on the null model. Most important for our present purposes here is that D_n asymptotically agrees with the $o(n)$ term in Theorem 2's bound. So, if we simulate data and evaluate D_n , then we can model D_n as a function of n to approximate its growth and, in turn, estimate the size of the $o(n)$ error term in Theorem 2.

We simulate 250 data sets for each value of n along the grid 2000, 4000, \dots , 30000 and, for each data set, evaluate the difference D_n . Define \bar{D}_n to be the average of D_n over the replicates, and then we model $\log \bar{D}_n$ as a linear function of $\log n$ —the coefficient associated with this $\log n$ term represents the polynomial order of $o(n)$, which should be much smaller than 1. We also fit a model that includes a $\log \log n$ term, but found that the coefficient on this term is not significantly different from zero. The specific mixture model under consideration here uses a Gaussian kernel $p_u(x) = \mathbf{N}(x | u, 1)$ mixed over the mean in $\mathbb{U} = [-5, 5]$. For the case when P^* is in the interior of the convex hull, we take the true mixing distribution to be $\text{Beta}(3, 5)$, centered and scaled to \mathbb{U} . For the boundary P^* case, we take the true mixing distribution to be a point mass at 0. These experiments were also carried with location-scale normal mixtures and mixtures of gamma kernels and the results were similar to those here.

Figure 4 shows our simulation-based estimates of the growth of the $o(n)$ error term as a function of n , with Panels (a) and (b) showing the interior- and boundary- P^* cases, respectively. In the interior case, for the linear model $\log \bar{D}_n \sim \log n$ stated using R syntax, the estimated coefficient on the $\log n$ term is 0.325 with standard error 0.002. This is in the ballpark of the conjectured value of 0.33 for this interior- P^* case. For the boundary- P^* case, the estimated coefficient is 0.545 with standard error 0.008, which is a little smaller than the conservative bound from [Martin and Tokdar \(2009\)](#).

As an aside, the near-2/3 lower bound on PR's weight sequence parameter γ appears to be an artifact of the proof technique employed in [Martin and Tokdar \(2009\)](#). One would expect that a smaller value of γ , e.g., close to the natural lower bound of 0.5, would be a better choice because the slower the weight sequence decays, the more opportunity an individual data point has to influence PR's estimator. So, intuitively, one would expect PR to perform better, to make more efficient use of the data, with smaller choices of the parameter γ . If true, then there would be a need for a different strategy for analyzing PR, one that could detect its improved performance with smaller weights. To check this intuition, and to see if we could get a smaller bound on the $o(n)$ term in Theorem 2,




Figure 4: Plots of \bar{D}_n versus n for two different simulation settings. In Panel (a), data are simulated from a P^* in the interior of the convex hull \mathcal{Q} whereas, in Panel (b), the data are simulated from P^* on the boundary of \mathcal{Q} .

we repeated the above experiment with $\gamma = 0.51$. Contrary to the above intuition, our findings revealed that the aforementioned $o(n)$ term is actually significantly larger for $\gamma = 0.51$ than it is for $\gamma = 0.67$. This suggests that there may be something fundamental about the near- $2/3$ bound in [Martin and Tokdar \(2009\)](#) after all.

F Testing log-concavity: multivariate case

One of the applications presented in the main paper was the use of e-processes, in particular, our proposed PRE-process, for testing log-concavity of the underlying density based on iid observations. There we focused on a comparison between our PRE-process-based test and the universal inference-based e-process test proposed by [Gangrade et al. \(2023\)](#). This latter e-process is very similar to the PRE-process: the only difference is that, instead of using PR to fit a nonparametric mixture density $\hat{q}_{X^{i-1}}(X_i)$ to each X_i based on X^{i-1} , it does so using a standard Gaussian kernel density estimator applied to X^{i-1} . The comparisons in the main paper focused on data of dimension $d = 1$. Our goal here is to do a similar comparison for $d > 1$.

First, we need to comment on the computation of the two e-processes when the dimension is greater than 1. Lets start with the denominators. As in the main paper, for dimension $d = 1$, the log-concave maximum likelihood can be found using the `logConDens` function from the R package `logcondens` ([Dümbgen and Rufibach 2011](#)); for $d \in \{2, 3, 4\}$ this can be done using the R package at <https://github.com/FabianRathke/fmlogcondens>. Next, we discuss computation of the e-process numerators in turn.

- For the kernel-based universal inference, we obtain the d -dimensional kernel density estimate using the `kde` function in the R package `ks` ([Duong 2024](#)); the same choice is made in [Gangrade et al. \(2023\)](#). This can accommodate dimension $d \leq 6$. We use a diagonal, data-driven bandwidth matrix as implemented in the `Hpi.diag` function in the same R package. Doing this sequentially, i.e., constructing a kernel

density estimator at X_i based on X^{i-1} for each $i \geq 1$, is quite expensive and, for this reason, [Gangrade et al. \(2023\)](#) propose a batching scheme. Here we are more interested in the quality of the e-processes than in the computational efficiency, so we did not implement their batched version.

- For PR, implementation is easy when the dimension is low. Here, however, for d -dimensional data, to fit a multivariate normal mixture, where mixing is over the mean vector and diagonal covariance matrix, then the dimension of the latent \mathbb{U} -space is $2d$. This means we quickly surpass cases where PR’s integration can be done using simple quadrature rules. It is for precisely these cases that we developed the PRticle filter strategy in [Dixit and Martin \(2023a\)](#), which goes roughly as follows. Start with a random sample $\{U_t : t = 1, \dots, T\}$ from PR’s initial guess $\hat{\Psi}_0$ of the mixing distribution supported on \mathbb{U} and, to each U_t , attach the weight $\pi_0(U_t) \equiv 1$. As described in [Dixit and Martin \(2023a\)](#), Sec. 3), the weights are sequentially updated as data points arrive so that the weighted samples after observing X^{i-1} can be used to approximate expectations with respect to the PR estimator $\hat{\Psi}_{i-1}$. In particular, they propose

$$\tilde{q}_{X^{i-1}}(X_i) := \frac{1}{T} \sum_{t=1}^T k(X_i | U_t) \pi_{i-1}(U_t)$$

as an approximation of PR’s $\hat{q}_{X^{i-1}}(X_i) = \int k(X_i | u) \Psi_{i-1}(du)$. Importantly, the weight updating scheme requires no integration. With this approximation of PR’s predictive density, it is straightforward to get the corresponding PRE-process.

To compare the two e-processes, we follow the strategy in Section [E](#) above. That is, we compare each of the e-processes to the oracle e-process that knows the true density p^* . In particular, we consider

$$D_n^{\text{PR}} = \log E_n^{\text{OR}} - \log E_n^{\text{PR}} \quad \text{and} \quad D_n^{\text{UI}} = \log E_n^{\text{OR}} - \log E_n^{\text{UI}},$$

where E_n^{PR} and E_n^{UI} are the PRE-process and kernel-based universal inference e-process, respectively. The advantage of this comparison is that it alleviates the computation of the log-concave maximum likelihood estimators and focuses specifically on the aspect that distinguishes the two e-processes, namely, their choice of numerator.

We focus our comparison on the case of $d = 3$. Since the kernel density estimation scheme can only handle $d \leq 6$, we consider the choice $d = 3$ to be “moderate dimension.” In this case, implementation of PR—fitting a Gaussian mixture model with diagonal covariance matrix—would require integration over a 6-dimensional space, which cannot be done using quadrature. Here, instead, we use the PRticle filter approximation described above with $T = 20,000$ initial particles. This can all be done up to at least dimension $d = 5$, but all the computations (both kernel and PR) become more expensive. Since the point we want to make is apparent with $d = 3$, this is where we focus our investigation.

Figure [5](#) shows the paths $n \mapsto D_n^\square$ for the two e-processes, $\square \in \{\text{PR}, \text{UI}\}$, under investigation over 50 data sets sampled iid from a true distribution P^* that is a two-component mixture of 3-dimensional Gaussian kernels with different means and different diagonal covariance matrices. This means that, in a certain sense, both e-processes are fitting a




Figure 5: Plots of D_n^{PR} (blue) and D_n^{UI} (pink) versus n for 50 data sets, along with the pointwise average.

“correctly-specified model” in their respective numerators, just with a different estimator: PR versus a kernel density estimator. Since D_n^{\square} measures the gap between the given e-process and the idealized oracle e-process, smaller values of D_n^{\square} indicate a smaller gap and, hence, a better approximation of that ideal. As is clear from the plot, D_n^{PR} tends to be significantly smaller than D_n^{UI} and far less variable across replications. The take-away message is that, while the increased dimension is sure to affect the quality of an e-process’s approximation of the oracle e-process, but the increased dimension has a more significant effect on the kernel-based e-process than it does on the PR-e-process implemented with the PRparticle filter. Moreover, although we did not implement the batch procedure suggested in [Gangrade et al. \(2023\)](#), it is worth mentioning that the PR-e-process’s denominator is much less expensive to compute than the kernel-based e-process’s, despite the fact that the latter is working with a relatively large number of particles. This is because the PR updates are genuinely recursive while a naive implementation of the kernel update, with data-driven bandwidth, requires a look-back at the full data.

G An alternative PR-like e-process

While there are a number of upsides to the use of a mixture model and PR in the denominator of the e-process construction, there are also downsides. One is that it is required to specify the mixture model’s kernel $\{p_u : u \in \mathbb{U}\}$. In some cases, this specification is trivial, i.e., one might just adopt a mixture of normals for the sake of flexibility, but in other cases it might not be. A natural question is, if all that one needs is the sequence of one-step-ahead predictive densities at the observations, then why is the mixture model formulation necessary? Can we not just work directly with densities on the sample space? [Hahn et al. \(2018\)](#) considered this question in the context of recursive Bayesian updating, and a proposal they put forward is the following. Their primary focus was on scalar

observations so that is the context we consider here. Like with PR, start with a sequence of weights (w_i) and a prior (univariate) predictive density q_0 with corresponding distribution function Q_0 . For a sequence of observations X_1, X_2, \dots , they suggest the following sequence of predictive densities

$$q_n(x) = (1 - w_n) q_{n-1}(x) + w_n q_{n-1}(x) c_\rho(Q_{n-1}(x), Q_{n-1}(X_n)), \quad x \in \mathbb{X}, \quad n \geq 1,$$

where c_ρ is the Gaussian copula density,

$$c_\rho(u, v) = \frac{\mathbf{N}_2(\Phi^{-1}(u), \Phi^{-1}(v) \mid 0, 1, \rho)}{\mathbf{N}(\Phi^{-1}(u) \mid 0, 1) \mathbf{N}(\Phi^{-1}(v) \mid 0, 1)},$$

with Φ the standard normal distribution function, and \mathbf{N}_2 in the numerator the bivariate normal density, with zero means and unit variances, depending on a correlation parameter $\rho \in [0, 1)$. Again, the advantage here compared to the original PR is that it does not require specification of a mixture model or the numerical integration to convert PR's mixing distribution estimator into a corresponding mixture density estimator. The recursion is written most naturally in terms of distribution—rather than density—functions, but then numerical differentiation is needed to get the density for the likelihood calculations that follow.

From the predictive updates, one can get a corresponding joint density via multiplication:

$$q^{\text{RBP}}(X^n) = \prod_{i=1}^n q_{i-1}(X_i),$$

where RBP indicates that this is motivated by recursive Bayes predictive updating. Given a null hypothesis \mathcal{P}_0 , one can define the corresponding PR-like e-process as

$$E_n^{\text{RBP}} = \frac{q^{\text{RBP}}(X^n)}{\sup_{p_0 \in \mathcal{P}_0} p_0(X^n)}.$$

This alternative e-process is interesting, and here are some points that we think warrant further investigation. First, there is currently no theory available that can be used to investigate the asymptotic growth rate of E_n^{RBP} , so can anything be said about this? Second, how does this perform empirically compared to the PRE-process and other available e-processes? Third, can this be extended naturally—and in a way that can be theoretically analyzed—to handle multivariate observations? A suggestion for handling bivariate data is given in [Hahn et al. \(2018\)](#), along with some brief numerical results, but no general formulation or theory.