

Optional Stopping with Bayes Factors: a categorization and extension of folklore results, with an application to invariant situations

Allard Hendriksen, Rianne de Heide, Peter Grünwald

July 25, 2018

Abstract

It is often claimed that Bayesian methods, in particular Bayes factor methods for hypothesis testing, can deal with optional stopping. We first give an overview, using only most elementary probability theory, of three different mathematical meanings that various authors give to this claim: stopping rule *independence*, posterior *calibration* and (semi-) *frequentist robustness to optional stopping*. We then prove theorems to the effect that — while their practical implications are sometimes debatable — these claims do indeed hold in a general measure-theoretic setting. The novelty here is that we allow for nonintegrable measures based on improper priors, which leads to particularly strong results for the practically important case of models satisfying a group invariance (such as location or scale). When equipped with the right Haar prior, calibration and semi-frequentist robustness to optional stopping hold uniformly irrespective of the value of the underlying nuisance parameter, as long as the stopping rule satisfies a certain intuitive property.

1 Introduction

In recent years, a surprising number of scientific results have failed to hold up to continued scrutiny. Part of this ‘replicability crisis’ may be caused by practices that ignore the assumptions of traditional (frequentist) statistical methods (John et al., 2012). One of these assumptions is that the experimental protocol should be completely determined upfront. In practice, researchers often adjust the protocol due to unforeseen circumstances or collect data until a point has been proven. This practice, which is referred to as *optional stopping* can cause true hypotheses to be wrongly rejected much more often than these statistical methods promise.

Bayes factor hypothesis testing has long been advocated as an alternative to traditional testing that can resolve several of its problems; in particular, it was claimed early on that Bayesian methods continue to be valid under optional stopping (Lindley, 1957, Edwards et al., 1963). In particular, the latter paper claims that (with Bayesian methods) “it is entirely appropriate to collect data until a point has been proven or disproven, or until the data collector runs out of time, money, or patience.”. In light of the replicability crisis, such claims have received much renewed interest (Wagenmakers, 2007, Rouder, 2014, Schönbrodt et al., 2017, Yu et al., 2013, Sanborn and Hills, 2013). But what do they mean mathematically? It turns out that different authors mean quite different things by ‘Bayesian methods handle optional stopping’; moreover, such claims are often shown to hold only in an informal sense, or in restricted contexts. Thus, the first goal of the present paper is to give a systematic overview and formalization of such claims in a simple, expository setting; the second goal is to extend the reach of such claims to more general settings, for which they have never been formally verified.

Overview and Most Important Result In *Section 2*, we give a systematic overview of what we identified to be the three main mathematical senses in which Bayes factor methods can handle optional stopping, which we call τ -independence, calibration, and (semi-) frequentist. We first do this in a setting chosen to be as simple as possible — finite sample spaces and strictly positive probabilities — allowing for straightforward statements and proofs of results.

In *Section 3*, we extend the statements and results to a much more general setting allowing for a wide range of sample spaces and measures, including measures based on *improper priors*. These are priors that are not integrable, thus not defining standard probability distributions over parameters, and as such they cause technical complications. Such priors are indispensable within the recently popularized, *default Bayes factors* for common hypothesis tests (Rouder et al., 2009, 2012, Jamil et al., 2016).

In *Section 4*, we provide stronger results for the case in which both models satisfy the same group invariance. Most (but not all) default Bayes factor settings concern such situations; prominent examples are Jeffreys’ Bayesian one- and two-sample t -tests, going back to (Jeffreys, 1961), in which the models are location and location-scale families, respectively. Many more examples are given by Berger and various collaborators in a sequence of papers (Berger et al., 1998, Dass and Berger, 2003, Bayarri et al., 2012, 2016) who give compelling arguments for using the (typically improper) *right Haar prior* on the nuisance parameters in such situations; for example, in Jeffreys’ one-sample t -test, one puts a right Haar prior on the variance. Haar priors and group invariant models were studied extensively by Eaton (1989), Andersson (1982), Wijsman (1990), whose results this paper depends on considerably. When nuisance parameters (shared by both H_0 and H_1) are of the right form and the right Haar prior is used, we can strengthen the results of *Section 3*: they now hold uniformly for all possible values of the nuisance parameters, rather than in the marginal, ‘on average’ sense we consider in *Section 3*; however — and this is our most important insight — we *cannot take arbitrary stopping rules* if we want to handle optional stopping in this strong sense: the stopping rules have to satisfy a certain intuitive condition, which will hold in many but not all practical cases: a rule such as ‘stop as soon as the Bayes factor is ≥ 20 ’ is allowed, but a rule (in the Jeffreys’ one-sample t -test) such as ‘stop as soon as $\sum x_i^2 \geq 20$ ’ is not.

The paper ends with an *Appendix* containing all longer mathematical proofs.

Remarks Our analysis is restricted to Bayesian testing and model selection using the Bayes factor method; we do not make any claims about other types of Bayesian inference. Some of the results we present were already known, at least in simple settings; we refer in each case to the first appearance in the literature that we are aware of. The main mathematical novelties in the paper are the results on optional stopping in the general case with improper priors and in the group invariance case. The main difficulties here are that, (a), for fixed sample sizes, at least with continuous-valued data, the Bayes factor usually has a distribution with full support, i.e. its density is strictly positive on the positive reals, whereas with variable stopping times, the support of its distribution may have ‘gaps’ at which its density is zero or very near 0; and (b), as indicated, we need certain restrictions on the stopping times in order for the results to be valid.

Finally, as an important caveat, we point out that the idea that Bayesian methods can handle optional stopping has also been criticized, for example by Yu et al. (2013), Sanborn and Hills (2013), and also by ourselves (De Heide and Grünwald, 2018). There are two main issues: first, in many practical situations, many Bayesian statisticians use priors that are *themselves* dependent on parts of the data and/or the sampling plan and stopping time. Examples are Jeffreys prior with the multinomial model and the Gunel-Dickey default priors for 2x2 contingency tables advocated by Jamil et al. (2016). With such priors, final results evidently depend on the stopping rule employed; none of the results below continue to hold for such priors. The second issue is that all mathematical theorems below are just that, mathematical theorems. For them to have implications for practice, one needs to make additional assumptions which sometimes may not be warranted. In particular the second sense of ‘handling

optional stopping’, calibration, relies on an analysis that holds under a Bayes marginal distribution, assigning probabilities that are really average (expected) probabilities with expectations taken over a prior. Yet many if not most priors used in practice (such as Cauchy priors over a location parameter) are of a ‘default’ or ‘pragmatic’ nature and are not really believed by the statistician, making the practical meaning of such an expectation questionable; [De Heide and Grünwald \(2018\)](#) discuss the issue at length.

2 The Simple Case

Consider a finite set \mathcal{X} and a sample space $\Omega := \mathcal{X}^T$ where T is some very large (but in this section, still finite) integer. One observes a *sample* $x^\tau \equiv x_1, \dots, x_\tau$, which is an initial segment of $x_1, \dots, x_T \in \mathcal{X}^T$. In the simplest case, $\tau = n$ is a sample size that is fixed in advance; but, more generally τ is a *stopping time* defined by some stopping rule (which may or may not be known to the data analyst), defined formally below.

We consider a hypothesis testing scenario where we wish to distinguish between a null hypothesis H_0 and an alternative hypothesis H_1 . Both H_0 and H_1 are sets of distributions on Ω , and they are each represented by unique probability distributions \bar{P}_0 and \bar{P}_1 respectively. Usually, these are taken to be Bayesian marginal distributions, defined as follows. First one writes, for both $k \in \{0, 1\}$, $H_k = \{P_{\theta|k} : \theta \in \Theta_k\}$ with ‘parameter spaces’ Θ_k ; one then defines or assumes some prior probability distributions π_0 and π_1 on Θ_0 and Θ_1 , respectively. The Bayesian marginal probability distributions are then the corresponding marginal distributions, i.e. for any set $A \subset \Omega$ they satisfy:

$$\bar{P}_0(A) = \int_{\Theta_0} P_{\theta|0}(A) d\pi_0(\theta) \quad ; \quad \bar{P}_1(A) = \int_{\Theta_1} P_{\theta|1}(A) d\pi_1(\theta). \quad (1)$$

For now we also further assume that for every $n \leq T$, every $x^n \in \mathcal{X}^n$, $\bar{P}_0(X^n = x^n) > 0$ and $\bar{P}_1(X^n = x^n) > 0$ (full support), where here as below we use random variable notation, $X^n = x^n$ denoting the event $\{x^n\} \subset \Omega$. We note that there exist approaches to testing and model choice such as testing by nonnegative martingales ([Shafer et al., 2011](#), [van der Pas and Grünwald, 2018](#)) and minimum description length ([Barron et al., 1998](#), [Grünwald, 2007](#)) in which the \bar{P}_0 and \bar{P}_1 may be defined in different (yet related) ways. Several of the results below extend to general \bar{P}_0 and \bar{P}_1 ; we return to this point at the end of the paper, in Section 5. In all cases, we further assume that we have determined an additional probability mass function π on $\{H_0, H_1\}$, indicating the prior probabilities of the hypotheses. The evidence in favour of H_1 relative to H_0 given data x^τ is now measured either by the *Bayes factor* or the *posterior odds*. We now give the standard definition of these quantities for the case that $\tau = n$, i.e., that the sample size is fixed in advance. First, noting that all conditioning below is on events of strictly positive probability, by Bayes’ theorem, we can write for any $A \subset \Omega$,

$$\frac{\pi(H_1 | A)}{\pi(H_0 | A)} = \frac{P(A | H_1)}{P(A | H_0)} \cdot \frac{\pi(H_1)}{\pi(H_0)}, \quad (2)$$

where here as in the remainder of the paper we use the symbol π to denote not just prior, but also posterior distributions on $\{H_0, H_1\}$. In the case that we observe x^n for fixed n , the event A is of the form $X^n = x^n$. Plugging this into (2), the left-hand side becomes the standard definition of *posterior odds*, and the first factor on the right is called the *Bayes factor*.

2.1 First Sense of Handling Optional Stopping: τ -Independence

Now, in reality we do not necessarily observe $X^n = x^n$ for fixed n but rather $X^\tau = x^\tau$ where τ is a stopping time that may itself depend on (past) data (and that in some cases may in fact be unknown

to us). This may be defined in terms of a *stopping rule* $f : \bigcup_{i \geq 0} \mathcal{X}^i \rightarrow \{\text{stop}, \text{continue}\}$. $\tau \equiv \tau(x^T)$ is then defined as the random variable which, for any sample x_1, \dots, x_T , outputs the smallest n such that $f(x_1, \dots, x_n) = \text{stop}$. For any given stopping time τ , any $1 \leq n \leq T$ and sequence of data $x^n = (x_1, \dots, x_n)$, we say that x^n is *compatible with* τ if it satisfies $X^n = x^n \Rightarrow \tau = n$. We let $\mathcal{X}^\tau \subset \bigcup_{i=1}^T \mathcal{X}^i$ be the set of all sequences compatible with τ .

Observations take the form $X^\tau = x^\tau$, which is equivalent to the event $X^n = x^n; \tau = n$ for some n and some $x^n \in \mathcal{X}^n$ which of necessity must be compatible with τ . We can thus instantiate (2) to

$$\frac{\pi(H_1 | X^n = x^n, \tau = n)}{\pi(H_0 | X^n = x^n, \tau = n)} = \frac{\bar{P}_1(X^n = x^n | \tau = n)}{\bar{P}_0(X^n = x^n | \tau = n)} \cdot \frac{\pi(H_1 | \tau = n)}{\pi(H_0 | \tau = n)} \quad (3)$$

where we note that, by definition $P_j(X^n = x^n | \tau = n) = P(X^n = x^n | H_j, \tau = n)$. Using Bayes' theorem again, the expression on the right can be further rewritten as:

$$\frac{\bar{P}_1(X^n = x^n | \tau = n)}{\bar{P}_0(X^n = x^n | \tau = n)} \cdot \frac{\pi(H_1 | \tau = n)}{\pi(H_0 | \tau = n)} = \frac{\bar{P}_1(X^n = x^n)/\bar{P}_1(\tau = n)}{\bar{P}_0(X^n = x^n)/\bar{P}_0(\tau = n)} \cdot \left(\frac{\bar{P}_1(\tau = n)}{\bar{P}_0(\tau = n)} \cdot \frac{\pi(H_1)}{\pi(H_0)} \right). \quad (4)$$

Combining (3) and (4) we get:

$$\frac{\overbrace{\pi(H_1 | X^n = x^n, \tau = n)}^{\gamma(x^n)}}{\overbrace{\pi(H_0 | X^n = x^n, \tau = n)}^{\beta(x^n)}} = \frac{\overbrace{\bar{P}_1(X^n = x^n)}^{\beta(x^n)}}{\bar{P}_0(X^n = x^n)} \cdot \frac{\pi(H_1)}{\pi(H_0)} \quad (5)$$

where we introduce the notation $\gamma(x^n)$ for the posterior odds and $\beta(x^n)$ for the Bayes factor based on sample x^n , calculated as if n were fixed in advance.

We see that the stopping rule plays no role in the expression on the right. Thus, we have shown that, for any two stopping times τ_1 and τ_2 that are both compatible with some observed x^n , the posterior odds one arrives at will be the same irrespective of whether x^n came to be observed because τ_1 was used or if x^n came to be observed because τ_2 was used. We say that the posterior odds do not depend on the stopping rule τ and call this property *τ -independence*. Incidentally, this also justifies that we write the posterior odds as $\gamma(x^n)$, a function of x^n alone, without referring to the stopping time τ .

The fact that the posterior odds given x^n do not depend on the stopping rule is the first (and still relatively weak) sense in which Bayesian methods handle optional stopping; it was perhaps first noted by Lindley (1957); another early source is Edwards et al. (1963). Lindley gave an (informal) proof in the context of specific parametric models; in Section 3.1 we show that the result indeed remains true for general σ -finite \bar{P}_0 and \bar{P}_1 . We note once again that the result only holds if the priors π_0 and π_1 themselves do not depend on τ , an assumption that is violated in many default Bayesian methods.

2.2 Second Sense of Handling Optional Stopping: Calibration

An alternative definition of handling optional stopping was introduced by Rouder (2014). Rouder calls $\gamma(x^n)$ the *nominal* posterior odds calculated from an obtained sample x^n , and defines the *observed posterior odds* as

$$\frac{\pi(H_1 | \gamma(x^n) = c)}{\pi(H_0 | \gamma(x^n) = c)}$$

as the posterior odds given the nominal odds. Rouder first notes that, at least if the sample size is fixed in advance to n , one expects these odds to be equal. For instance, if an obtained sample yields nominal posterior odds of 3-to-1 in favor of the alternative hypothesis, then it must be 3 times

as likely that the sample was generated by the alternative probability measure. In the terminology of De Heide and Grünwald (2018), Bayes is *calibrated* for a fixed sample size n . Rouder then goes on to note that, if n is determined by an arbitrary stopping time τ (based for example on optional stopping), then the odds will still be equal — in this sense, Bayesian testing is well-behaved in the calibration sense irrespective of the stopping rule/time. Formally, the requirement that the nominal and observed posterior odds be equal leads us to define the *calibration hypothesis*, which postulates that $c = \frac{P(H_1|\gamma=c)}{P(H_0|\gamma=c)}$ holds for any $c > 0$ that has non-zero probability. For simplicity, for now we only consider the case with equal prior odds for H_0 and H_1 so that $\gamma(x^n) = \beta(x^n)$. Then the calibration hypothesis says that, for arbitrary stopping time τ , for every c such that $\beta(x^\tau) = c$ for some $x^\tau \in \mathcal{X}^\tau$, one has

$$c = \frac{P(\beta(x^\tau) = c | H_1)}{P(\beta(x^\tau) = c | H_0)}. \quad (6)$$

In the present simple setting, this hypothesis is easily shown to hold, because we can write:

$$\frac{P(\beta(X^\tau) = c | H_1)}{P(\beta(X^\tau) = c | H_0)} = \frac{\sum_{y \in \mathcal{X}^\tau: \beta(y)=c} P(\{y\} | H_1)}{\sum_{y \in \mathcal{X}^\tau: \beta(y)=c} P(\{y\} | H_0)} = \frac{\sum_{y \in \mathcal{X}^\tau: \beta(y)=c} c P(\{y\} | H_0)}{\sum_{y \in \mathcal{X}^\tau: \beta(y)=c} P(\{y\} | H_0)} = c.$$

Rouder noticed that the calibration hypothesis should hold as a mathematical theorem, without giving an explicit proof; he demonstrated it by computer simulation in a simple parametric setting. Deng et al. (2016) gave a proof for a somewhat more extended setting yet still with proper priors. In Section 3.2 we show that a version of the calibration hypothesis continues to hold for general measures based on improper priors, and in Section 4.4 we extend this further to strong calibration for group invariance settings as discussed below.

We note that this result, too, relies on the priors themselves not depending on the stopping time, an assumption which is violated in several standard default Bayes factor settings. We also note that, if one thinks of one's priors in a default sense — they are practical but not necessarily fully believed — then the practical implications of calibration are limited, as shown experimentally by De Heide and Grünwald (2018). One would really like a stronger form of calibration in which (6) holds under a whole range of distributions in H_0 and H_1 , rather than in terms of \bar{P}_0 and \bar{P}_1 which average over a prior that perhaps does not reflect one's beliefs fully. For the case that H_1 and H_2 share a nuisance parameter g taking values in some set G , one can define this *strong calibration hypothesis* as stating that, for all c with $\beta(x^\tau) = c$ for some $x^\tau \in \mathcal{X}^\tau$, all $g \in G$,

$$c = \frac{P(\beta(x^\tau) = c | H_1, g)}{P(\beta(x^\tau) = c | H_0, g)}. \quad (7)$$

where β is still defined as above; in particular, when calculating β one does not condition on the parameter having the value g , but when assessing its likelihood as in (7) one does. De Heide and Grünwald (2018) show that the strong calibration hypothesis certainly does *not* hold for general parameters, but they also show by simulations that it does hold in the practically important case with group invariance and right Haar priors. In Section 4.4 we show that in such cases, one can indeed prove that a version of (7) holds.

2.3 Third Sense of Handling Optional Stopping: (Semi-) Frequentist

In classical, Neyman-Pearson style null hypothesis testing, the main concern is limiting the false positive rate of a hypothesis test. If this false positive rate is bounded above by some $\alpha > 0$, then a null hypothesis significance test (NHST) is said to have *significance level* α , and if the significance level is independent of the stopping rule used, we say that the test is *robust under frequentist optional stopping*.

Definition 1. A function $S : \bigcup_{i=m}^T \mathcal{X}^i \rightarrow \{0, 1\}$ is said to be a frequentist sequential test with significance level α and minimal sample size m that is robust under optional stopping relative to H_0 if for all $P \in H_0$

$$P(\exists n \in \{m+1, \dots, T\} : S(X^n) = 1) \leq \alpha,$$

i.e. the probability that there is an n at which $S(X^n) = 1$ ('the test rejects H_0 when given sample X^n ') is bounded by α .

In our present setting, we can take $m = 0$ (larger m become important in Section 3.3), so n runs from 1 to T and it is easy to show that, for any $0 \leq \alpha \leq 1$, we have

$$\bar{P}_0 \left(\exists n \in \{1, \dots, T\} : \frac{1}{\beta(x^n)} \leq \alpha \right) \leq \alpha. \quad (8)$$

Proof. For any fixed α and any sequence $x^T = x_1, \dots, x_T$, let $\tau(x^T)$ be the smallest n such that, for the initial segment x^n of x^T , $\beta(x^n) \geq 1/\alpha$ (if no such n exists we set $\tau(x^T) = T$). Then τ is a stopping time, X^τ is a random variable, and the probability in (8) is equal to the \bar{P}_0 -probability that $\beta(X^\tau) \geq 1/\alpha$ which by Markov's inequality is bounded by α . \square

It follows that, if H_0 is a singleton, then the sequential test S that rejects H_0 (outputs $S(X^n) = 1$) whenever $\beta(x^n) \geq 1/\alpha$ is a frequentist sequential test with significance level α that is robust under optional stopping.

The fact that Bayes factor testing with singleton H_0 handles optional stopping in this frequentist way was noted by Edwards et al. (1963) and also emphasized by Good (1991), among many others. If H_0 is not a singleton, then (8) still holds, so the Bayes factor still handles optional stopping in a mixed frequentist (Type I-error) and Bayesian (marginalizing over prior within H_0) sense. While from a frequentist perspective, one may not consider this to be fully satisfactory, as argued by Bayarri et al. (2016) this semi-frequentist sense of optional stopping, when applied to H_1 rather than H_0 , sometimes does correspond to what many frequentists consider acceptable in practice.

Yet, in the practically important group invariance case, we can once again show that, if all parameters in H_0 are right Haar, then the Bayes factor is truly robust to optional stopping in the above frequentist sense, i.e. (8) will hold for all $P \in H_0$ and not just 'on average'. While this is hinted at in several papers (e.g. (Bayarri et al., 2016, Dass and Berger, 2003)) it seems to never have been proven formally; we provide a proof in Section 4.5.

3 The General Case

Let (Ω, \mathcal{F}) be a measurable space. Fix some $m \geq 0$ and consider a sequence of functions X_{m+1}, X_{m+2}, \dots on Ω so that each X_n , $n > m$ takes values in some fixed set ('outcome space') \mathcal{X} with associated σ -algebra Σ . When working with proper priors we invariably take $m = 0$ and then we define $X^n := (X_1, X_2, \dots, X_n)$ and we let $\Sigma^{(n)}$ be the n -fold product algebra of Σ . When working with improper priors it turns out to be useful (more explanation further below) to take $m > 0$ and define an *initial sample* random variable $\langle X^{(m)} \rangle$ on Ω , taking values in some set $\langle \mathcal{X}^m \rangle \subseteq \mathcal{X}^m$ with associated σ -algebra $\langle \Sigma^{(m)} \rangle$. In that case we set, for $n \geq m$, $\langle \mathcal{X}^n \rangle = \{x^n = (x_1, \dots, x_n) \in \mathcal{X}^n : x^m = (x_1, \dots, x_m) \in \langle \mathcal{X}^m \rangle\}$, and $X^n := (\langle X^{(m)} \rangle, X_{m+1}, X_{m+2}, \dots, X_n)$ and we let $\Sigma^{(n)}$ be $\langle \Sigma^{(m)} \rangle \times \prod_{j=m+1}^n \Sigma$. In either case, we let \mathcal{F}_n be the σ -algebra (relative to Ω) generated by $(X^n, \Sigma^{(n)})$. Then $(\mathcal{F}_n)_{n=m, m+1, \dots}$ is a filtration relative to \mathcal{F} and if we equip (Ω, \mathcal{F}) with a distribution P then $\langle X^{(m)} \rangle, X_{m+1}, X_{m+2}, \dots$ becomes a random process adapted to \mathcal{F} . A *stopping time* is now generalized to be a function

$\tau : \Omega \rightarrow \{m+1, m+2, \dots\} \cup \{\infty\}$ such that for each $n > m$, the event $\{\tau = n\}$ is \mathcal{F}_n -measurable; note that we only consider stopping after m initial outcomes. Again, for a given stopping time τ and sequence of data $x^n = (x_1, \dots, x_n)$, we say that x^n is compatible with τ if it satisfies $X^n = x^n \Rightarrow \tau = n$, i.e. $\{\omega \in \Omega : X^n(\omega) = x^n\} \subset \{\omega \in \Omega : \tau(\omega) = n\}$.

H_0 and H_1 are now sets of probability distributions on (Ω, \mathcal{F}) . Again one writes $H_j = \{P_{\theta|j} : \theta \in \Theta_j\}$ where now the parameter sets Θ_j (which, however, could itself be infinite-dimensional) are themselves equipped with suitable σ -algebras $\Sigma(\Theta_j)$.

We will still represent both H_0 and H_1 by unique measures \bar{P}_0 and \bar{P}_1 respectively, which we now allow to be based on (1) with improper priors π_0 and π_1 that may be infinite measures; as a result \bar{P}_0 and \bar{P}_1 are positive real measures that may themselves be infinite. We also allow \mathcal{X} to be a general (in particular uncountable) set. Both nonintegrability and uncountability cause complications, but these can be overcome if suitable Radon-Nikodym derivatives exist. To ensure this, we will assume that for all $n \geq \max\{m, 1\}$, for all $k \in \{0, 1\}$ and $\theta \in \Theta_k$, $P_{\theta|k}^{(n)}$, $\bar{P}_0^{(n)}$ and $\bar{P}_1^{(n)}$ are all mutually absolutely continuous and that the measures $\bar{P}_1^{(n)}$ and $\bar{P}_0^{(n)}$ are σ -finite. Then there also exists a measure ρ on (Ω, \mathcal{F}) such that, for all such n , $\bar{P}_1^{(n)}$, $\bar{P}_0^{(n)}$ and $\rho^{(n)}$ are all mutually absolutely continuous: we can simply take $\rho^{(n)} = \bar{P}_0^{(n)}$, but in practice, it is often possible and convenient to take ρ such that $\rho^{(n)}$ is the Lebesgue measure on \mathbb{R}^n , which is why explicitly introduce ρ here.

The absolute continuity conditions guarantee that all required Radon-Nikodym derivatives exist. Finally, we assume that the measures $\bar{P}_0^{(n)}$ and $\bar{P}_1^{(n)}$ are proper probability measures for all $n > m$. This final requirement is the reason why we sometimes need to consider $m > 0$ and nonstandard sample spaces $\langle \mathcal{X}^n \rangle$ in the first place: in practice, one usually starts with the standard setting of a (Ω, \mathcal{F}) where $m = 0$ and all X_i have the same status. In all practical situations with improper priors π_0 and/or π_1 that we know of, there is a smallest finite j and a set $\mathcal{X}^\circ \subset \mathcal{X}^j$ that has measure 0 under all probability distributions in $H_0 \cup H_1$, such that, restricted to the sample space $\mathcal{X}^j \setminus \mathcal{X}^\circ$, the measures $\bar{P}_1^{(j)}$ and $\bar{P}_0^{(j)}$ are σ -finite and mutually absolutely continuous, and the posteriors $\pi_k(\Theta_k | x^j)$ (as defined in the standard manner in (11) below) are proper probability measures. One then sets m to equal this j , and sets $\langle \mathcal{X}^m \rangle := \mathcal{X}^m \setminus \mathcal{X}^\circ$, and the required properness will be guaranteed. Our initial sample $\langle X^{(m)} \rangle$ is a variation of what is called (for example, by Bayarri et al. (2012)) a *minimal sample*. Yet, the sample size of a standard minimal sample is itself a random quantity; by restricting \mathcal{X}^m to $\langle \mathcal{X}^m \rangle$, we can take its sample size m to be constant rather than random, which will greatly simplify the treatment of optional stopping with group invariance; see Example 1 and 2 below.

We henceforth refer to the setting now defined (with m and initial space $\langle \mathcal{X}^m \rangle$ satisfying the requirements above) as the *general case*.

We need an analogue of (5) for this general case. If \bar{P}_0 and \bar{P}_1 are probability measures, then there is still a standard definition of conditional probability distributions $P(H | \mathcal{A})$ in terms of conditional expectation for any given σ -algebra \mathcal{A} ; based on this, we can derive the required analogue in two steps. First, we consider the case that $\tau \equiv n$ for some $n > m$; we know in advance that we observe X^n for a fixed n : the appropriate \mathcal{A} is then \mathcal{F}_n , $\pi(H | \mathcal{A})(\omega)$ is determined by $X^n(\omega)$ hence can be written as $\pi(H | X^n)$, and a straightforward calculation gives that

$$\frac{\pi(H_1 | X^n = x^n)}{\pi(H_0 | X^n = x^n)} = \left(\left(\frac{d\bar{P}_1^{(n)}/d\rho^{(n)}}{d\bar{P}_0^{(n)}/d\rho^{(n)}} \right) (x^n) \right) \cdot \frac{\pi(H_1)}{\pi(H_0)} \quad (9)$$

where $(d\bar{P}_1^{(n)}/d\rho^{(n)})$ and $(d\bar{P}_0^{(n)}/d\rho^{(n)})$ are versions of the Radon-Nikodym derivatives defined relative to $\rho^{(n)}$. The second step is now to follow exactly the same steps as in the derivation of (5), replacing $\beta(X^n)$ by (9) wherever appropriate (we omit the details). This yields, for any n such that $\rho(\tau = n) > 0$,

and for $\rho^{(n)}$ -almost every x^n that is compatible with τ ,

$$\frac{\overbrace{\pi(H_1 | x^n)}^{\gamma_n}}{\pi(H_0 | x^n)} = \frac{\pi(H_1 | X^n = x^n, \tau = n)}{\pi(H_0 | X^n = x^n, \tau = n)} = \overbrace{\left(\frac{d\bar{P}_1^{(n)}/d\rho^{(n)}}{d\bar{P}_0^{(n)}/d\rho^{(n)}} \right) (x^n)}^{\beta_n} \cdot \frac{\pi(H_1)}{\pi(H_0)}, \quad (10)$$

where here as below, for $n \geq m$, we abbreviate $\pi(H_k | X^n = x^n)$ to $\pi(H_k | x^n)$.

The above expression for the posterior is valid if \bar{P}_0 and \bar{P}_1 are probability measures; we will simply take it as the *definition* of the Bayes factor for the general case. Again this coincides with standard usage for the improper prior case. In particular, let us define the conditional posteriors and Bayes factors given $\langle X^{(m)} \rangle = x^m$ in the standard manner, by the formal application of Bayes' rule, for $k = 0, 1$ and measurable $\Theta'_k \subset \Theta_k$ and \mathcal{F} -measurable A ,

$$\pi_k(\Theta'_k | x^m) := \frac{\int_{\Theta'_k} \frac{dP_{\theta|k}^{(m)}}{d\rho^{(m)}}(x^m) d\pi_k(\theta)}{\int_{\Theta_k} \frac{dP_{\theta|k}^{(m)}}{d\rho^{(m)}}(x^m) d\pi_k(\theta)} \quad (11)$$

$$\bar{P}_k(A | x^m) := \bar{P}_k(A | \langle X^{(m)} \rangle = x^m) := \int_{\Theta_k} P_{\theta|k}(A | \langle X^{(m)} \rangle = x^m) d\pi_k(\theta | x^m), \quad (12)$$

where $P_{\theta|k}(A | \langle X^{(m)} \rangle = x^m)$ is defined as the value that (a version of) the conditional probability $P_{\theta|k}(A | \mathcal{F}_m)$ takes when $\langle X^{(m)} \rangle = x^m$, and is thus defined up to a set of $\rho^{(m)}$ -measure 0.

With these definitions, it is straightforward to derive we have the following *coherence property*, which automatically holds if the priors are proper, and which in combination with (10) expresses that first updating on x^m and then on x_{m+1}, \dots, x_n has the same result as updating based on the full x_1, \dots, x_n at once:

$$\frac{\pi(H_1 | X^n = x^n, \tau = n)}{\pi(H_0 | X^n = x^n, \tau = n)} = \left(\frac{d\bar{P}_1^{(n)}(\cdot | x^m)}{d\bar{P}_0^{(n)}(\cdot | x^m)}(x^n) \right) \cdot \frac{\pi(H_1 | x^m)}{\pi(H_0 | x^m)}. \quad (13)$$

3.1 τ -independence, general case

The general version of the claim that the posterior odds do not depend on the specific stopping rule that was used is now immediate, since the expression (10) for the Bayes factor does not depend on the stopping time τ .

3.2 Calibration, general case

We will now show that the calibration hypothesis continues to hold in our general setting. From here onward, we make the further reasonable assumption that for every $x^m \in \langle \mathcal{X}^m \rangle$, $\bar{P}_0(\tau = \infty | x^m) = \bar{P}_1(\tau = \infty | x^m) = 0$ (the stopping time is almost surely finite), and we define $\mathcal{T}_\tau := \{n \in \mathbb{N}_{>0} : \bar{P}_0(\tau = n) > 0\}$.

To prepare further, let $\{B_j : j \in \mathcal{T}_\tau\}$ be any collection of positive random variables such that for each $j \in \mathcal{T}_\tau$, B_j is \mathcal{F}_j -measurable. We can define the *stopped* random variable B_τ as

$$B_\tau := \sum_{j=0}^{\infty} \mathbb{1}_{\{\tau=j\}} B_j, \quad (14)$$

where we note that, under this definition, B_τ is well-defined even if $\mathbf{E}_{\bar{P}_0}[\tau] = \infty$.

We can define the induced measures on the positive real line under the null and alternative hypothesis for any probability measure P on (Ω, \mathcal{F}) :

$$P^{[B_\tau]} : \mathcal{B}(\mathbb{R}_{>0}) \rightarrow [0, 1] : A \mapsto P(B_\tau^{-1}(A)). \quad (15)$$

where $\mathcal{B}(\mathbb{R}_{>0})$ denotes the Borel σ -algebra of $\mathbb{R}_{>0}$. Note that, when we refer to $P^{[B_n]}$, this is identical to $P^{[B_\tau]}$ for the stopping time τ which on all of Ω stops at n . The following lemma is crucial for passing from fixed-sample size to stopping-rule based results.

Lemma 1. *Let \mathcal{T}_τ and $\{B_n : n \in \mathcal{T}_\tau\}$ be as above. Consider two probability measures P_0 and P_1 on (Ω, \mathcal{F}) . Suppose that for all $n \in \mathcal{T}_\tau$, the following fixed-sample size calibration property holds:*

$$\text{for some fixed } c > 0, \quad P_0^{[B_n]} \text{-almost all } b : \frac{P_1(\tau = n)}{P_0(\tau = n)} \cdot \frac{dP_1^{[B_n]}(\cdot | \tau = n)}{dP_0^{[B_n]}(\cdot | \tau = n)}(b) = c \cdot b. \quad (16)$$

Then we have

$$\text{for } P_0^{[B_\tau]} \text{-almost all } b : \frac{dP_1^{[B_\tau]}}{dP_0^{[B_\tau]}}(b) = c \cdot b. \quad (17)$$

The proof is in Appendix A.

In this subsection we apply this lemma to the measures $\bar{P}_k(\cdot | x^m)$ for arbitrary fixed $x^m \in \langle \mathcal{X}^m \rangle$, with their induced measures $\bar{P}_0^{[\gamma_\tau]}(\cdot | x^m)$, $\bar{P}_1^{[\gamma_\tau]}(\cdot | x^m)$ for the *stopped posterior odds* γ_τ . Formally, the posterior odds γ_n as defined in (10) constitute a random variable for each n , and, under our mutual absolute continuity assumption for \bar{P}_0 and \bar{P}_1 , γ_n can be directly written as $\frac{d\bar{P}_1^{(n)}}{d\bar{P}_0^{(n)}} \cdot \pi(H_1)/\pi(H_0)$. Since, by definition, the measures $\bar{P}_k(\cdot | x^m)$ are probability measures, the Radon-Nikodym derivatives in (16) and (17) are well-defined.

Lemma 2. *We have for all $x^m \in \langle \mathcal{X}^m \rangle$, all $n > m$:*

$$\text{for } \bar{P}_0^{[\gamma_n]}(\cdot | x^m) \text{-almost all } b : \frac{\bar{P}_1^{[\gamma_n]}(\tau = n | x^m)}{\bar{P}_0^{[\gamma_n]}(\tau = n | x^m)} \cdot \frac{d\bar{P}_1^{[\gamma_n]}(\cdot | x^m)}{d\bar{P}_0^{[\gamma_n]}(\cdot | x^m)}(b) = \frac{\pi(H_0 | x^m)}{\pi(H_1 | x^m)} \cdot b. \quad (18)$$

Combining the two lemmas now immediately gives (19) below, and combining further with (13) and (10) gives (20):

Corollary 3. *In the setting considered above, we have for all $x^m \in \langle \mathcal{X}^m \rangle$:*

$$\text{for } \bar{P}_0^{[\gamma_\tau]}(\cdot | x^m) \text{-almost all } b : \frac{\pi(H_1 | x^m)}{\pi(H_0 | x^m)} \cdot \frac{d\bar{P}_1^{[\gamma_\tau]}(\cdot | x^m)}{d\bar{P}_0^{[\gamma_\tau]}(\cdot | x^m)}(b) = b, \quad (19)$$

and also

$$\text{for } \bar{P}_0^{[\gamma_\tau]}(\cdot | x^m) \text{-almost all } b : \frac{\pi(H_1)}{\pi(H_0)} \cdot \frac{d\bar{P}_1^{[\gamma_\tau]}}{d\bar{P}_0^{[\gamma_\tau]}}(b) = b, \quad (20)$$

In words, the posterior odds remain calibrated under any stopping rule τ which stops almost surely at times $m < \tau < \infty$.

For discrete and strictly positive measures with prior odds $\pi(H_1)/\pi(H_0) = 1$, we always have $m = 0$, and (19) is equivalent to (6). Note that $\bar{P}_0^{[\gamma_\tau]} x^m$ -a.s. in (19) is equivalent to $\bar{P}_1^{[\gamma_\tau]} x^m$ -a.s. because the two measures are assumed mutually absolutely continuous.

3.3 (Semi-) Frequentist Optional Stopping

In this section we consider our general setting as in the beginning of Section 3.2, i.e. with the added assumption that the stopping time is a.s. finite, and with $\mathcal{T}_\tau := \{j \in \mathbb{N}_{>0} : \bar{P}_0(\tau = j) > 0\}$. From here onward we shall further simplify slightly by assuming that the prior on H_0 and H_1 is equal, so that the posterior odds (given x^n and possibly $\tau = n$) equal the Bayes factor β_n .

Consider any initial sample $x^m \in \langle \mathcal{X}^m \rangle$ and let $\bar{P}_0 | x^m$ and $\bar{P}_1 | x^m$ be the conditional Bayes marginal distributions as defined in (12). We first note that, by Markov's inequality, for any nonnegative random variable Z on Ω with, for all $x^m \in \langle \mathcal{X}^m \rangle$, $\mathbf{E}_{\bar{P}_0 | x^m}[Z] \leq 1$, we must have, for $0 \leq \alpha \leq 1$, $\bar{P}_0(Z^{-1} \leq \alpha | x^m) \leq \mathbf{E}_{\bar{P}_0 | x^m}[\beta_\tau]/\alpha^{-1} \leq \alpha$.

Proposition 4. *Let τ be any stopping rule satisfying our requirements. The stopped Bayes factor β_τ as defined by (14) (with β_j in the role of B_j) is a random variable that satisfies, for all $x^m \in \langle \mathcal{X}^m \rangle$, $\mathbf{E}_{\bar{P}_0 | x^m}[\beta_\tau] \leq 1$, so that, by the reasoning above, $\bar{P}_0(\frac{1}{\beta_\tau} \leq \alpha | x^m) \leq \alpha$.*

Proof. The following implications are all immediate:

$$\begin{aligned} \mathbf{E}_{\bar{P}_0 | x^m}[\beta_\tau] &= \mathbf{E}_{\bar{P}_0 | x^m} \left[\sum_{n: \bar{P}_0(\tau=n | x^m) > 0} \mathbb{1}_{\{\tau=n\}} \beta_n \right] \\ &= \sum_{n: \bar{P}_0(\tau=n | x^m) > 0} \mathbb{1}_{\{\tau=n\}} \bar{P}_0(\tau = n | x^m) \cdot \mathbf{E}_{\bar{P}_0 | x^m, \tau=n}[\beta_n] \\ &= \sum_{n: \bar{P}_0(\tau=n | x^m) > 0} \mathbb{1}_{\{\tau=n\}} \bar{P}_0(\tau = n | x^m) \cdot \int \frac{d\bar{P}_1^{(n)}(\cdot | x^m, \tau = n)}{d\bar{P}_0^{(n)}(X^n | x^m, \tau = n)} d(\bar{P}_0 | x^m, \tau) = 1. \end{aligned}$$

□

The desired result now follows by plugging in a particular stopping rule: let $S : \bigcup_{i=m}^{\infty} \mathcal{X}^i \rightarrow \{0, 1\}$ be the frequentist sequential test defined by setting, for all $n > m$, $x^n \in \langle \mathcal{X}^n \rangle$: $S(x^n) = 1$ iff $\beta_n \geq 1/\alpha$.

Corollary 5. *Let $t^* \in \{m+1, m+2, \dots\} \cup \{\infty\}$ be the smallest $t^* > m$ for which $\beta_{t^*}^{-1} \leq \alpha$. Then for arbitrarily large T , when applied to the stopping rule $\tau := \min\{T, t^*\}$, we find that*

$$\bar{P}_0(\exists n \in \{m+1, \dots, T\} : S(X^n) = 1 | x^m) = \bar{P}_0(\exists n \in \{m+1, \dots, T\} : \beta_n^{-1} \leq \alpha | x^m) \leq \alpha.$$

The corollary implies that the test S is robust under optional stopping in the frequentist sense relative to \mathcal{H}_0 (Definition 1). Note that, just as in the simple case, the setting is really just ‘semi-frequentist’ whenever H_0 is not a singleton.

4 Optional stopping with group invariance

Whenever the null hypothesis is composite, the previous results only hold under the marginal distribution \bar{P}_0 or, in the case of improper priors, under $\bar{P}_0(\cdot | X^m = x^m)$. When a group structure can be imposed on the outcome space and (a subset of the) parameters that is joint to H_0 and H_1 , stronger results can be derived for calibration and frequentist optional stopping. Invariably, such parameters function as *nuisance parameters* and our results are obtained if we equip them with the so-called *right Haar prior* which is usually improper. Below we show how we then obtain results that simultaneously hold for *all* values of the nuisance parameters. Such cases include many standard testing scenarios such as the (Bayesian variations) of the t -test, as illustrated in the examples below. Note though that our

results do not apply to settings with improper priors for which no group structure exists; for example, if $P_{\theta|0}$ expresses that X_1, X_2, \dots are i.i.d. $\text{Poisson}(\theta)$, then from an objective Bayes or MDL point of view it makes sense to adopt Jeffreys' prior for the Poisson model; this prior is improper, allows initial sample size $m = 1$, but does not allow for a group structure. For such a prior we can only use the marginal results Corollary 3 and Corollary 5.

4.1 Background for fixed sample sizes

Here we prepare for our results by providing some general background on invariant priors for Bayes factors with fixed sample size n on models with nuisance parameters that admit a group structure, introducing the right Haar measure, the corresponding Bayes marginals, and (maximal) invariants. We use these results in Section 4.2 to derive Lemma 7, which gives us a strong version of calibration for fixed n . The setting is extended to variable stopping times in Section 4.3, and then Lemma 7 is used in this extended setting to obtain our strong optional stopping results in Section 4.4 and 4.5.

For now, we assume a sample space $\langle \mathcal{X}^n \rangle$ that is locally compact and Hausdorff, and that is a subset of some product space \mathcal{X}^n where \mathcal{X} is itself locally compact and Hausdorff. This requirement is met, for example, when $\mathcal{X} = \mathbb{R}$ and $\langle \mathcal{X}^n \rangle = \mathcal{X}^n$. In practice, the space $\langle \mathcal{X}^n \rangle$ is invariably a subset of \mathcal{X}^n that excludes some 'singular' outcomes that have measure 0 under all hypotheses involved. We associate $\langle \mathcal{X}^n \rangle$ with its Borel σ -algebra which we denote as \mathcal{F}_n . Observations are denoted by the random vector $X^n = (X_1, \dots, X_n) \in \langle \mathcal{X}^n \rangle$. We thus consider outcomes of fixed sample size, denoting these as $x^n \in \langle \mathcal{X}^n \rangle$, returning to the case with stopping times in Section 4.4 and 4.5.

We start with some group-theoretical preliminaries; for more details, see e.g. (Eaton, 1989, Wijsman, 1990, Andersson, 1982).

Definition 2 (Eaton (1989), Definition 2.1). *Let G be a group of measurable one-to-one transformations on \mathcal{X} with identity e . Let Y be a set. A function $F : Y \times G \rightarrow Y$ satisfying*

1. $F(y, e) = y, y \in Y$
2. $F(y, g_1 g_2) = F(F(y, g_1), g_2), g_1, g_2 \in G, y \in Y$

specifies G acting on the right of Y .

In practice, F is omitted: we will write $y \cdot g$ for a group element g acting on the right of $y \in Y$. For a subset $A \subseteq Y$, we write $A \cdot g := \{a \cdot g \mid a \in A\}$. From now on we let G be a locally compact group G that acts topologically and properly¹ on the right of $\langle \mathcal{X}^n \rangle$.

Let $P_{0,e}$ and $P_{1,e}$ (notation to become clear below) be two arbitrary probability distributions on $\langle \mathcal{X}^n \rangle$ that are mutually absolutely continuous. We will now generate hypothesis classes H_0 and H_1 , both sets of distributions on $\langle \mathcal{X}^n \rangle$ with parameter space G , starting from $P_{0,e}$ and $P_{1,e}$. The group action of G on $\langle \mathcal{X}^n \rangle$ induces a group action on these measures defined by

$$P_{k,g}(A) := (P_{k,e} \cdot g)(A) := P_{k,e}(A \cdot g^{-1}) = \int \mathbb{1}_{\{A\}}(x \cdot g) P_{k,e}(dx) \quad (21)$$

for any set $A \in \mathcal{F}_n$, $k = 0, 1$. When applied to $A = \langle \mathcal{X}^n \rangle$, we get $P_{k,g}(A) = 1$, for all $g \in G$, whence we have created two sets of probability measures parameterized by g , i.e.,

$$H_0 := \{P_{0,g} \mid g \in G\} \quad ; \quad H_1 := \{P_{1,g} \mid g \in G\}. \quad (22)$$

¹A group acts properly on a set Y if the mapping $\psi : Y \times G \mapsto Y \times Y$ defined by $\psi(y, g) = (y \cdot g, y)$ is a proper mapping, i.e. the inverse image of ψ of each compact set in $Y \times Y$ is a compact set in $Y \times G$. (Eaton (1989), Definition 5.1)

In this context, $g \in G$, can typically be viewed as nuisance parameter, i.e. a parameter that is not directly of interest, but needs to be accounted for in the analysis. This is illustrated in Example 1 and Example 2 below. The examples also illustrate how to extend this setting to cases where there are more parameters than just $g \in G$ in either H_0 or H_1 . We extend the whole setup to our general setting with non-fixed n in Section 4.4.

We use the right Haar measure for G as a prior to define the Bayes marginals:

Definition 3 (Conway (2013), Example 1.11). *Let G be a locally compact topological group. Then the right invariant Haar measure (in short: right Haar measure) for G is a Borel measure ν satisfying*

1. $\nu(A) > 0$ for every nonempty open set $A \subseteq G$,
2. $\nu(K) < \infty$ for every compact set $K \subseteq G$,
3. $\nu(A \cdot g) = \nu(A)$ for every $g \in G$ and every measurable $A \subseteq G$.

The Bayes marginals are then defined to be:

$$\bar{P}_k(A) = \int_G \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{A\}} dP_{k,g} \nu(dg) \quad (23)$$

for $k = 0, 1$ and $A \in \mathcal{F}_n$. Typically, the right Haar measure is improper so that the Bayes marginals \bar{P}_k are not integrable. Yet, in all cases of interest, they are (a) still σ -finite, and, (b), \bar{P}_0, \bar{P}_1 and all distributions in $H_k, k = 0, 1, g \in G$ are mutually absolutely continuous; we will henceforth assume that (a) and (b) are the case.

Example 1. Consider the one-sample t-test as described by Rouder et al. (2009), going back to (Jeffreys, 1961). The test considers normally distributed data with unknown standard deviation. The test is meant to answer the question whether the data has mean $\mu = 0$ (the null hypothesis) or some other mean (the alternative hypothesis). Below, all densities are taken with respect to Lebesgue measure. Following (Rouder et al., 2009), a Cauchy prior density, denoted by $\pi_\delta(\delta)$, is placed on the effect size $\delta = \mu/\sigma$. The unknown standard deviation is a nuisance parameter and is equipped with the right Haar prior with density $\pi_\sigma(\sigma) = \frac{1}{\sigma}$ under both hypotheses. Abbreviating, for general measure P on $\langle \mathcal{X}^n \rangle$, $(dP/d\lambda)$ (the density of distribution P relative to Lebesgue measure on \mathbb{R}^n) to p , this gives the following marginal densities on n outcomes:

$$\begin{aligned} p_{0,\sigma}(x^n) &= \frac{1}{(2\pi\sigma^2)^{n/2}} \cdot \exp\left(-\frac{1}{2\pi\sigma^2} \sum_{i=1}^n x_i^2\right) \\ p_{1,\sigma,\delta}(x^n) &= \frac{1}{(2\pi\sigma^2)^{n/2}} \cdot \exp\left(-\frac{n}{2} \left[\left(\frac{\bar{x}}{\sigma} - \delta\right)^2 + \left(\frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2}{\sigma^2}\right) \right]\right) \quad \left(\text{with } \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i\right) \\ p_{1,\sigma}(x^n) &= \int p_{1,\sigma,\delta}(x^n) \pi_\delta(\delta) d\delta \\ \bar{p}_0(x^n) &= \int_0^\infty p_{0,\sigma}(x^n) \pi_\sigma(\sigma) d\sigma, \\ \bar{p}_1(x^n) &= \int_0^\infty \int_{-\infty}^\infty p_{1,\sigma,\delta}(x^n) \pi_\delta(\delta) \pi_\sigma(\sigma) d\delta d\sigma = \int_0^\infty p_{1,\sigma}(x^n) \pi_\sigma(\sigma) d\sigma. \end{aligned}$$

Normally, this is viewed as a test between $H_0 = \{P_{0,\sigma} : \sigma \in \mathbb{R}_{>0}\}$ and $H'_1 = \{P_{1,\sigma,\delta} : \sigma \in \mathbb{R}_{>0}, \delta \in \mathbb{R}\}$, but we can obviously also view it as test between H_0 and $H_1 = \{P_{1,\sigma}\}$, i.e. we integrate out the parameter δ . Then H_0 and H_1 will be in the form (22) needed to state our results. In all standard invariant settings, one can reduce H_0 and H_1 to this form in the same way, by integrating out all other

parameters, which is possible as long we equip those parameters with proper priors, which is invariably done in practice.

The group for the nuisance parameter σ is the group of scale transformations $G = \{c : c \in \mathbb{R}_{>0}\}$. We thus let the sample space be $\langle \mathcal{X}^n \rangle = \mathbb{R}^n \setminus \{0\}^n$: we remove the measure-zero set $\{0\}^n$, such that the group action is proper on the sample space. The group action is defined by $x^n \cdot c = cx^n$ for $x^n \in \langle \mathcal{X}^n \rangle, c \in G$. Take a $g \in G$. The measures $P_{k,g}$ defined by (21) are given by their densities $p_{0,\sigma}$ and $p_{1,\sigma}$ as defined above, with σ replaced by g .

We now turn to the main ingredient that will be needed to obtain results on optional stopping: the quotient σ -algebra.

Definition 4 (Eaton (1989), Chapter 2). *A group G acting on the right of a set Y induces an equivalence relation: $y_1 \sim y_2$ if and only if there exists $g \in G$ such that $y_1 = y_2 \cdot g$. This equivalence relation partitions the space in orbits: $O_y = \{y \cdot g \mid g \in G\}$, the collection of which is called the quotient space Y/G . There exists a map, the natural projection, from Y to the quotient space which is defined by $\varphi_Y : Y \rightarrow Y/G : y \mapsto \{y \cdot g \mid g \in G\}$, and which we use to define the quotient σ -algebra*

$$\mathcal{G}_n = \{\varphi_{\langle \mathcal{X}^n \rangle}^{-1}(\varphi_{\langle \mathcal{X}^n \rangle}(A)) \mid A \in \mathcal{F}_n\}. \quad (24)$$

Definition 5 (Eaton (1989), Chapter 2). *A random element U_n on $\langle \mathcal{X}^n \rangle$ is invariant if for all $g \in G, x^n \in \langle \mathcal{X}^n \rangle, U_n(x^n) = U_n(x^n \cdot g)$. The random element U_n is maximal invariant if U_n is invariant and for all $y^n \in \langle \mathcal{X}^n \rangle, U_n(x^n) = U_n(y^n)$ implies $x^n = y^n \cdot g$ for some $g \in G$.*

Thus, U_n is maximal invariant if and only if U_n is constant on each orbit, and takes different values on different orbits; $\varphi_{\langle \mathcal{X}^n \rangle}$ is thus an example of a maximal invariant. Note that any maximal invariant is \mathcal{G}_n -measurable. The importance of this quotient σ -algebra \mathcal{G}_n is the following evident fact:

Proposition 6. *For fixed $k \in \{0, 1\}$, every invariant U_n has the same distribution under all $P_{k,g}, g \in G$.*

Chapter 2 of (Eaton, 1989) provides several methods and examples how to construct a concrete maximal invariant, including the first two given below. Since β_n is invariant under the group action of G (see below), β_n is an example of an invariant, although not necessarily of a maximal invariant.

Example 1 (continued) Consider the setting of the one-sample t-test as described above in Example 1. A maximal invariant for $x^n \in \langle \mathcal{X}^n \rangle$ is $U_n(x^n) = (x_1/|x_1|, x_2/|x_1|, \dots, x_n/|x_1|)$.

Example 2. A second example, with a group invariance structure on two parameters, is the setting of the two-sample t-test with the right Haar prior (which coincides here with Jeffreys' prior) $\pi(\mu, \sigma) = 1/\sigma$ (see Rouder et al. (2009) for details): the group is $G = \{(a, b) : a > 0, b \in \mathbb{R}\}$. Let the sample space be $\langle \mathcal{X}^n \rangle = \mathbb{R}^n \setminus \text{span}(e_n)$, where e_n denotes a vector of ones of length n (this is to exclude the measure-zero line for which the $s(x^n)$ is zero), and define the group action by $x^n \cdot (a, b) = ax^n + be_n$ for $x^n \in \langle \mathcal{X}^n \rangle$. Then (Eaton (1989), Example 2.15) a maximal invariant for $x^n \in \langle \mathcal{X}^n \rangle$ is $U_n(x^n) = (x^n - \bar{x}e_n)/s(x^n)$, where \bar{x} is the sample mean and $s(x^n) = (\sum_{i=1}^n (x_i - \bar{x})^2)^{1/2}$.

However, we can also construct a maximal invariant similar to the one in Example 1, which gives a special status to an initial sample:

$$U_n(X^n) = \left(\frac{X_2 - X_1}{|X_2 - X_1|}, \frac{X_3 - X_1}{|X_2 - X_1|}, \dots, \frac{X_n - X_1}{|X_2 - X_1|} \right), \quad n \geq 2.$$

4.2 Relatively Invariant Measures and Calibration for Fixed n

Let U_n be a maximal invariant, taking values in measurable space $(\mathcal{U}_n, \mathcal{G}_n)$. Although we have given more concrete examples above, it follows from the results of [Andersson \(1982\)](#) that, in case we do not know how to construct a U_n , we can always take $U_n = \varphi_{\langle \mathcal{X}^n \rangle}$. Since we assume mutual absolute continuity, the Radon-Nikodym derivative $\frac{dP_{1,g}^{[U_n]}}{dP_{0,g}^{[U_n]}}$ must exist and we can apply the following theorem (note it is here that the use of *right* Haar measure is crucial; a different result holds for the left Haar measure):²

Theorem ([Berger et al., 1998, Theorem 2.1](#)) Under our previous definitions of and assumptions on G , $P_{k,g}$, \bar{P}_k let $\beta(x^n) := \bar{P}_1(x^n)/\bar{P}_0(x^n)$ be the Bayes factor based on x^n . Let U_n be a maximal invariant as above, with (adopting the notation of (15)) marginal measures $P_{k,g}^{[U_n]}$, for $k = 0, 1$ and $g \in G$. There exists a version of the Radon-Nikodym derivative such that we have for all $g \in G$, all $x^n \in \langle \mathcal{X}^n \rangle$,

$$\frac{dP_{1,g}^{[U_n]}}{dP_{0,g}^{[U_n]}}(U_n(x^n)) = \beta(x^n). \quad (25)$$

As a first consequence of the theorem above, we note (as did [Berger et al. \(1998\)](#)) that the Bayes factor $\beta_n := \beta(X^N)$ is \mathcal{G}_n -measurable (it is constant on orbits), and thus *it has the same distribution under $P_{0,g}$ and $P_{1,g}$ for all $g \in G$* . The theorem also implies the following crucial lemma:

Lemma 7. [Strong Calibration for Fixed n] *Under the assumptions of the theorem above, let U_n be a maximal invariant and let V_n be a \mathcal{G}_n -measurable binary random variable with $P_{0,g}(V_n = 1) > 0, P_{1,g}(V_n = 1) > 0$. Adopting the notation of (15), we can choose the Radon-Nikodym derivative $dP_{1,g}^{[\beta_n]}(\cdot | V_n = 1)/dP_{0,g}^{[\beta_n]}(\cdot | V_n = 1)$ so that we have, for all $x^n \in \langle \mathcal{X}^n \rangle$:*

$$\frac{P_{1,g}(V_n = 1)}{P_{0,g}(V_n = 1)} \cdot \frac{dP_{1,g}^{[\beta_n]}(\cdot | V_n = 1)}{dP_{0,g}^{[\beta_n]}(\cdot | V_n = 1)}(\beta_n(x^n)) = \beta_n(x^n), \quad (26)$$

where for the special case with $P_{k,g}(V_n = 1) = 1$, we get $\frac{dP_{1,g}^{[\beta_n]}}{dP_{0,g}^{[\beta_n]}}(\beta_n(x^n)) = \beta_n(x^n)$.

4.3 Extending to Our General Setting with Non-Fixed Sample Sizes

We start with the same setting as above: a group G on sample space $\langle \mathcal{X}^n \rangle \subset \mathcal{X}^n$ that acts topologically and properly on the right of $\langle \mathcal{X}^n \rangle$; two distributions $P_{0,e}$ and $P_{1,e}$ on $(\langle \mathcal{X}^n \rangle, \mathcal{F}_n)$ that are used to generate H_0 and H_1 , and Bayes marginal measures based on the right Haar measure \bar{P}_0 and \bar{P}_1 , which are both σ -finite. We now denote H_k as $H_k^{(n)}$, $P_{k,e}$ as $P_{k,e}^{(n)}$ and \bar{P}_k as $\bar{P}_k^{(n)}$, all $P \in H_0^{(n)} \cup H_1^{(n)}$ are mutually absolutely continuous.

We now extend this setting to our general random process setting as specified in the beginning of [Section 3.2](#) by further assuming that, for the same group G , for some $m > 0$, the above setting is defined for each $n \geq m$. To connect the $H_k^{(n)}$ for all these n , we further assume that there exists a subset $\langle \mathcal{X}^m \rangle \subset \mathcal{X}^m$ that has measure 1 under $P_{k,e}^{(n)}$ (and hence under all $P_{g,e}^{(n)}$) such that for all $n \geq m$:

²This theorem requires that there exists some relatively invariant measure μ on $\langle \mathcal{X}^n \rangle$ such that for $k = 0, 1, g \in G$, the $P_{k,g}$ all have a density relative to μ . Since the Bayes marginal \bar{P}_0 based on the right Haar prior is easily seen to be such a relatively invariant measure, the conditions for the theorem apply.

1. We can write $\langle \mathcal{X}^n \rangle = \{x^n \in \mathcal{X}^n : (x_1, \dots, x_m) \in \langle \mathcal{X}^m \rangle\}$.
2. For all $x^n \in \langle \mathcal{X}^n \rangle$, the posterior $\nu \mid x^n$ based on the right Haar measure ν is proper.
3. The probability measures $P_{k,e}^{(n)}$ and $P_{k,e}^{(n+1)}$ satisfy Kolmogorov's compatibility condition for a random process.
4. The group action \cdot on the measures $P_{k,e}^{(n)}$ and $P_{k,e}^{(n+1)}$ is compatible, i.e. for every $n > 0$, for every $A \in \mathcal{F}_n$, every $g \in G$, $k \in \{0, 1\}$, we have $P_{k,g}^{(n+1)}(A) = P_{k,g}^{(n)}(A)$.

Requirement 4. simply imposes the condition that the group action considered is the same for all $n \in \mathbb{N}$. As a consequence of 3. and 4., the probability measures $P_{k,g}^{(n)}$ and $P_{k,g}^{(n+1)}$ satisfy Kolmogorov's compatibility condition for all $g \in G$, $k \in \{0, 1\}$ which means that there exists a probability measure $P_{k,g}$ on (Ω, \mathcal{F}) (under which $\langle X^{(m)} \rangle, X_{m+1}, X_{m+2}, \dots$ is a random process), defined as in the beginning of Section 3, whose marginals for $n \geq m$ coincide with $P_{k,g}^{(n)}$, and there exist measures \bar{P}_0 and \bar{P}_1 on (Ω, \mathcal{F}) whose marginals for $n \geq m$ coincide with $\bar{P}_0^{(n)}$ and $\bar{P}_1^{(n)}$. We have thus defined a set H_0 and H_1 of hypotheses on (Ω, \mathcal{F}) and the corresponding Bayes marginals \bar{P}_0 and \bar{P}_1 and are back in our general setting. It is easily verified that the 1- and 2-sample Bayesian t-tests both satisfy all these assumptions: in Example 1, take $m = 1$ and $\langle \mathcal{X}^m \rangle = \mathbb{R} \setminus \{0\}$; in Example 4.1, take $m = 2$ and $\langle \mathcal{X}^m \rangle = \mathbb{R}^2 \setminus \{(a, a) : a \in \mathbb{R}\}$. The conditions can also be verified for the variety of examples considered by Berger et al. (1998) and Bayarri et al. (2012). In fact, our initial sample $x^m \in \langle \mathcal{X}^m \rangle$ is a variation of what they call a *minimal sample*; by excluding 'singular' outcomes from \mathcal{X}^m to ensure that the group acts properly on $\langle \mathcal{X}^m \rangle$, we can guarantee that the initial sample is of fixed size; the size of the minimal sample can be larger, on a set of measure 0 under all $P \in H_0 \cup H_1$, e.g. if, in Example 4.1, $X_1 = X_2$. We chose to ensure a fixed size m since it makes the extension to random processes considerably easier.

Just as in the case with a fixed n , if we start with large (or even nonparametric) hypotheses $H'_k = \{P_{\theta'|k} : \theta' \in \Theta'_k\}$ which we want to equip with a nuisance parameter g , we can still view this as an instance of the present setting by first taking a proper prior density $d\pi(\theta')$ on θ' , setting $P_{k,e}$ equal to the corresponding Bayes marginal, i.e.

$$P_{k,e} := \int P_{\theta'} \pi(\theta') d\pi(\theta'), \quad (27)$$

and defining $P_{k,g}^{(n)}$ for each n (and hence $P_{k,g}$) via (21). We can then view H_k equivalently as $\{P_{k,g} : g \in G\}$, with \bar{P}_k the Bayes marginal based on the right Haar prior (and Θ_k in (1) is equal to G); or as

$$H_k = \{P_{(\theta',g)|k} : \theta' \in \Theta'_k, g \in G\}, \quad (\text{with } \Theta_k = (\Theta'_k, G)) \quad (28)$$

with now \bar{P}_k the Bayes marginal based on the product prior of $\pi(\theta')$ and the right Haar prior. This was illustrated (with θ' equal to the effect size) for the 1-sample Bayesian t -test in Example 1.

4.4 Strong Calibration

Consider the setting, definitions and assumptions of the previous subsection, with the additional assumptions and definitions made in the beginning of Section 3.3, in particular the assumption of a.s. finite stopping time and equal prior odds. We will now show a strong calibration theorem for the Bayes factors $\beta_n = (d\bar{P}_0^{(n)})/(d\bar{P}_1^{(n)})(X^n)$ defined in terms of the Bayes marginals \bar{P}_0 and \bar{P}_1 with the right Haar prior. Thus β_τ is defined as in (14) with β in the role of B .

Theorem 8 (Strong calibration under optional stopping). *Let τ be a stopping time satisfying our requirements, such that additionally, for each $n > m$, the event $\{\tau = n\}$ is \mathcal{G}_n -measurable. Then, adopting the notation of (15), for all $g \in G$, for $P_{0,g}^{[\beta_\tau]}$ -almost every $b > 0$, we have:*

$$\frac{dP_{1,g}^{[\beta_\tau]}}{dP_{0,g}^{[\beta_\tau]}}(b) = b.$$

That means that the posterior odds remain calibrated under every stopping rule τ adapted to the quotient space filtration $\mathcal{G}_m, \mathcal{G}_{m+1}, \dots$, under all $P_{0,g}$.

Proof. Fix some $g \in G$. We simply first apply Lemma 7 with $V_n = \mathbb{1}_{\{\tau=n\}}$, which gives that the premise (16) of Lemma 1 holds with $c = 1$ and β_n in the role of B_n (it is here that we need that τ_n is \mathcal{G}_n -measurable, otherwise we could not apply Lemma 7 with the required definition of V_n). We can now use Lemma 1 with $P_{0,g}$ in the role of P_0 to reach the desired conclusion for the chosen g . Since this works for all $g \in G$, the result follows. \square

Example 1, Continued: Admissible and Inadmissible Stopping Rules We obtain strong calibration for the one-sample t-test with respect to the nuisance parameter σ (see Example 1 above) when the stopping rule is adapted to the quotient filtration $\mathcal{G}_m, \mathcal{G}_{m+1}, \dots$. Under each $P_{k,g} \in H_k$, the Bayes factors $\beta_m, \beta_{m+1}, \dots$ define a random process on Ω such that each β_n is \mathcal{G}_n -measurable. This means that a stopping time defined in terms of rule such as ‘stop at the smallest t at which $\beta_t > 20$ or $t = 10^6$ ’ is allowed in the result above. If the stopping rule is a function of a sequence of maximal invariants, like $x_1/|x_1|, x_2/|x_2|, \dots$, it is adapted to the \mathcal{G}_n as well and we can again apply the result above. On the other hand, this requirement is violated, for example, by a stopping rule that stops when $\sum_{i=1}^j (x_i)^2$ exceeds some fixed value, since such a stopping rule explicitly depends on the scale of the sampled data.

4.5 Frequentist optional stopping

The special case of the following result for the one-sample Bayesian t -test was proven in the master’s thesis (Hendriksen, 2017). Here we extend the result to general group invariances.

Theorem 9 (Frequentist optional stopping for composite null hypotheses with group invariance). *Under the same conditions as in Section 4.4, let τ be a stopping time such that, for each $n > m$, the event $\{\tau = n\}$ is \mathcal{G}_n -measurable. Then, adopting the notation of (15), for all $g \in G$, the stopped Bayes factor satisfies $\mathbf{E}_{P_{0,g}[\beta_\tau]} = \int_{\mathbb{R}_{>0}} c dP_{0,g}^{[\beta_\tau]}(c) = 1$, so that, by the reasoning above Proposition 4, we have for all $g \in G$: $P_{0,g}(\frac{1}{\beta_\tau} \leq \alpha) \leq \alpha$.*

Proof. [of Theorem 9] We have

$$\int_{\mathbb{R}_{>0}} c dP_{0,g}^{[\beta_\tau]}(c) = \int_{\mathbb{R}_{>0}} \frac{dP_{1,g}^{[\beta_\tau]}}{dP_{0,g}^{[\beta_\tau]}}(c) dP_{0,g}^{[\beta_\tau]}(c) = \int_{\mathbb{R}_{>0}} dP_{1,g}^{[\beta_\tau]}(c) = 1.$$

where the first equality follows directly from Theorem 8 and the final equality follows because $P_{1,g}$ is a probability measure, integrating to 1. \square

Analogously to Corollary 5, the desired result now follows by plugging in a particular stopping rule: let $S : \bigcup_{i=m}^{\infty} \mathcal{X}^i \rightarrow \{0, 1\}$ be the frequentist sequential test defined by setting, for all $n > m$, $x^n \in \langle \mathcal{X}^n \rangle$: $S(x^n) = 1$ iff $\beta_n \geq 1/\alpha$.

Corollary 10. Let $t^* \in \{m + 1, m + 2, \dots\} \cup \{\infty\}$ be the smallest $t^* > m$ for which $\beta_{t^*}^{-1} \leq \alpha$. Then for arbitrarily large T , when applied to the stopping rule $\tau := \min\{T, t^*\}$, we find that for all $g \in G$:

$$P_{0,g}(\exists n \in \{m + 1, \dots, T\} : S(X^n) = 1 \mid x^m) = P_{0,g}(\exists n \in \{m + 1, \dots, T\} : \beta_n^{-1} \leq \alpha \mid x^m) \leq \alpha.$$

The corollary implies that the test S is robust under optional stopping in the frequentist sense relative to \mathcal{H}_0 (Definition 1).

Example 1 (continued) When we choose a stopping rule that is $(\mathcal{G}_m, \mathcal{G}_{m+1}, \dots)$ -measurable, the hypothesis test is robust under (semi-) frequentist optional stopping. This holds for example, for the one- and two-sample t-test (Rouder et al., 2009), Bayesian ANOVA (Rouder et al., 2012), and Bayesian linear regression (Liang et al., 2008); again, for stopping rules that are not $(\mathcal{G}_m, \mathcal{G}_{m+1}, \dots)$ -measurable, robustness under frequentist optional stopping can be violated.

5 Concluding Remarks

We have identified three types of ‘handling optional stopping’: τ -independence, calibration and semi-frequentist. We extended the corresponding definitions and results to general sample spaces with potentially improper priors. For the special case of models H_0 and H_1 sharing a nuisance parameter with a group invariant structure, we showed stronger versions of calibration and semi-frequentist robustness to optional stopping. As a final remark, it is worth noting that — as is immediate from the proofs — all our group-invariance results continue to hold in the setting with H_k as in (28), and the definition of the Bayes marginal $P_{k,e}$ relative to θ' as in (27) replaced by a probability measure on (Ω, \mathcal{F}) that is not necessarily of the Bayes marginal form. The results work for any probability measure; in particular one can take the alternatives for the Bayes marginal with proper prior that are considered in the minimum description length and sequential prediction literature (Barron et al., 1998, Grünwald, 2007) under the name of *universal distribution* relative to $\{P_{\theta'} : \theta' \in \Theta'\}$; examples include the prequential, normalized maximum likelihood or ‘switch’ distributions considered by van der Pas and Grünwald (2018).

Acknowledgments

We are grateful to Wouter Koolen, for extremely useful conversations that helped with the math, and to Jeff Rouder, for providing inspiration and insights that sparked off this research.

References

- S. Andersson. Distributions of maximal invariants using quotient measures. *The Annals of Statistics*, 10(3):955–961, 1982. ISSN 00905364. URL <http://www.jstor.org/stable/2240918>.
- A. Barron, J. Rissanen, and B. Yu. The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6):2743–2760, 1998. Special Commemorative Issue: Information Theory: 1948-1998.
- M. Bayarri, D. J. Benjamin, J. O. Berger, and T. M. Sellke. Rejection odds and rejection ratios: A proposal for statistical practice in testing hypotheses. *Journal of Mathematical Psychology*, 72: 90–103, 2016.

- M. J. Bayarri, J. O. Berger, A. Forte, G. García-Donato, et al. Criteria for Bayesian model choice with application to variable selection. *The Annals of statistics*, 40(3):1550–1577, 2012.
- J. O. Berger, L. R. Pericchi, and J. A. Varshavsky. Bayes factors and marginal distributions in invariant situations. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 307–321, 1998.
- J. B. Conway. *A course in functional analysis*, volume 96. Springer Science & Business Media, 2013.
- S. C. Dass and J. O. Berger. Unified conditional frequentist and Bayesian testing of composite hypotheses. *Scandinavian Journal of Statistics*, 30(1):193–210, Mar 2003. ISSN 1467-9469. doi: 10.1111/1467-9469.00326. URL <http://dx.doi.org/10.1111/1467-9469.00326>.
- A. Deng, J. Lu, and S. Chen. Continuous monitoring of A/B tests without pain: Optional stopping in Bayesian testing. In *Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on*, pages 243–252. IEEE, 2016.
- M. L. Eaton. Group invariance applications in statistics. *Regional Conference Series in Probability and Statistics*, 1:i–133, 1989. ISSN 19355912. URL <http://www.jstor.org/stable/4153172>.
- W. Edwards, H. Lindman, and L. J. Savage. Bayesian statistical inference for psychological research. *Psychological Review*, 70(3):193–242, 1963. ISSN 0033-295X. doi: 10.1037/h0044139. URL <http://dx.doi.org/10.1037/h0044139>.
- I. Good. C383. A comment concerning optional stopping. *Journal of Statistical Computation and Simulation*, 39(3):191–192, 1991.
- P. Grünwald. *The Minimum Description Length Principle*. MIT Press, Cambridge, MA, 2007.
- R. de Heide and P. Grünwald. Why optional stopping is a problem for Bayesians. *arXiv preprint arXiv:1708.08278*, 2018.
- A. A. Hendriksen. Betting as an alternative to p -values. Master’s thesis, Leiden University, Dept. of Mathematics, 2017.
- T. Jamil, A. Ly, R. Morey, J. Love, M. Marsman, and E.-J. Wagenmakers. Default “Gunel and Dickey” Bayes factors for contingency tables. *Behavior Research Methods*, pages 1–15, 2016.
- H. Jeffreys. *Theory of Probability*. Oxford, Oxford, England, 1961.
- L. K. John, G. Loewenstein, and D. Prelec. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*, 2012.
- F. Liang, L. Paulo, G. Molina, M. Clyde, and J. Berger. Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103(481):410–423, 2008.
- D. V. Lindley. A statistical paradox. *Biometrika*, 44(1/2):187, Jun 1957. ISSN 0006-3444. doi: 10.2307/2333251. URL <http://dx.doi.org/10.2307/2333251>.
- J. Rouder, R. Morey, P. Speckman, and J. Province. Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5):356–374, 2012.
- J. N. Rouder. Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review*, 21(2):301–308, Mar 2014. ISSN 1531-5320. doi: 10.3758/s13423-014-0595-4. URL <http://dx.doi.org/10.3758/s13423-014-0595-4>.
- J. N. Rouder, P. L. Speckman, D. Sun, R. D. Morey, and G. Iverson. Bayesian t -tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2):225–237, Apr 2009. ISSN 1531-5320. doi: 10.3758/pbr.16.2.225. URL <http://dx.doi.org/10.3758/pbr.16.2.225>.

- A. N. Sanborn and T. T. Hills. The frequentist implications of optional stopping on Bayesian hypothesis tests. *Psychonomic Bulletin & Review*, 21(2):283–300, Oct 2013. ISSN 1531-5320. doi: 10.3758/s13423-013-0518-9. URL <http://dx.doi.org/10.3758/s13423-013-0518-9>.
- F. D. Schönbrodt, E.-J. Wagenmakers, M. Zehetleitner, and M. Perugini. Sequential hypothesis testing with Bayes factors: Efficiently testing mean differences. *Psychological Methods*, 22(2):322–339, 2017. ISSN 1082-989X. doi: 10.1037/met0000061. URL <http://dx.doi.org/10.1037/met0000061>.
- G. Shafer, A. Shen, N. Vereshchagin, and V. Vovk. Test martingales, Bayes factors and p-values. *Statistical Science*, pages 84–101, 2011.
- S. van der Pas and P. Grünwald. Almost the best of three worlds: Risk, consistency and optional stopping for the switch criterion in nested model selection. *Statistica Sinica*, 28(1):229–255, 2018.
- E.-J. Wagenmakers. A practical solution to the pervasive problems of p-values. *Psychonomic Bulletin & Review*, 14(5):779–804, Oct 2007. ISSN 1531-5320. doi: 10.3758/bf03194105. URL <http://dx.doi.org/10.3758/BF03194105>.
- R. A. Wijsman. *Invariant measures on groups and their use in statistics*. Institute of Mathematical Statistics, 1990. ISBN 9780940600195. URL <http://www.worldcat.org/oclc/24066028?referer=xd>.
- E. C. Yu, A. M. Sprenger, R. P. Thomas, and M. R. Dougherty. When decision heuristics and science collide. *Psychonomic Bulletin & Review*, 21(2):268–282, Sep 2013. ISSN 1531-5320. doi: 10.3758/s13423-013-0495-z. URL <http://dx.doi.org/10.3758/s13423-013-0495-z>.

A Proofs Omitted from Main Text

Proof. [of Lemma 1] Let $A \subset \mathbb{R}_{>0}$ be any Borel measurable set. In the equations below, the sum and integral can be swapped due to the monotone convergence theorem and the fact that B_τ is a positive

function.

$$\begin{aligned}
\int_A dP_1^{[B_\tau]} &= \int_\Omega \mathbb{1}_{\{B_\tau \in A\}} dP_1^{[B_\tau]} \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_\tau \in A\}} \mathbb{1}_{\{\tau=n\}} dP_1^{(n)} \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_n \in A\}} \mathbb{1}_{\{\tau=n\}} dP_1^{(n)}(\cdot | \tau = n) \cdot P_1^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_n \in A\}} dP_1^{(n)}(\cdot | \tau = n) \cdot P_1^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{r>0} \mathbb{1}_{\{r \in A\}} dP_1^{[B_n]}(\cdot | \tau = n) \cdot P_1^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{r>0} \mathbb{1}_{\{r \in A\}} \frac{dP_1^{[B_n]}(\cdot | \tau = n)}{dP_0^{[B_n]}(\cdot | \tau = n)}(r) dP_0^{[B_n]}(\cdot | \tau = n) \cdot P_1^{(n)}(\tau = n) \\
&\stackrel{(*)}{=} \sum_{n=0}^{\infty} \int_{r>0} \mathbb{1}_{\{r \in A\}} \frac{P_0^{(n)}(\tau = n)}{P_1^{(n)}(\tau = n)} \cdot r dP_0^{[B_n]}(\cdot | \tau = n) \cdot P_1^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{r>0} \mathbb{1}_{\{r \in A\}} r dP_0^{[B_n]}(\cdot | \tau = n) \cdot P_0^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_n \in A\}} \cdot B_n dP_0^{(n)}(\cdot | \tau = n) P_0^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_n \in A\}} \mathbb{1}_{\{\tau=n\}} \cdot B_n dP_0^{(n)}(\cdot | \tau = n) P_0^{(n)}(\tau = n) \\
&= \sum_{n=0}^{\infty} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{B_\tau \in A\}} \mathbb{1}_{\{\tau=n\}} \cdot B_n dP_0^{(n)} \\
&= \int_\Omega \mathbb{1}_{\{B_\tau \in A\}} \left(\sum_{n=0}^{\infty} \mathbb{1}_{\{\tau=n\}} B_n \right) dP_0 = \int_A t P_0^{[B_\tau]}(dt),
\end{aligned}$$

where (*) follows because of our fixed n -calibration assumption.

We have shown that the function g defined by $g(t) = t$ is the Radon-Nikodym derivative $\frac{dP_1^{[B_\tau]}}{dP_0^{[B_\tau]}}$. \square

Proof. [of Lemma 2] Let A be any Borel subset of $\mathbb{R}_{>0}$. We have:

$$\begin{aligned}
\int_A d\bar{P}_1^{[\gamma_n]}(\cdot | x^m, \tau = n) &= \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{\gamma_n \in A\}} d\bar{P}_n^{(1)}(\cdot | x^m, \tau = n) \\
&= \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{\gamma_n \in A\}} \left(\frac{d\bar{P}_n^{(1)}(\cdot | x^m, \tau = n)}{d\bar{P}_n^{(0)}(\cdot | x^m, \tau = n)} \right) d\bar{P}_n^{(0)}(\cdot | x^m, \tau = n) \\
&\stackrel{(*)}{=} \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{\gamma_n \in A\}} \gamma_n \cdot \left(\frac{\pi(H_0 | x^m, \tau = n)}{\pi(H_1 | x^m, \tau = n)} \right) d\bar{P}_n^{(0)}(\cdot | x^m, \tau = n) \\
&= \int_{\langle \mathcal{X}^n \rangle} \mathbb{1}_{\{\gamma_n \in A\}} \gamma_n \left(\frac{\bar{P}_0(\tau = n | x^m) \pi(H_0)}{\bar{P}_1(\tau = n | x^m) \pi(H_1)} \right) d\bar{P}_n^{(0)}(\cdot | x^m, \tau = n) \\
&= \left(\frac{\bar{P}_0(\tau = n | x^m) \pi(H_0 | x^m)}{\bar{P}_1(\tau = n | x^m) \pi(H_1 | x^m)} \right) \cdot \int_A \gamma_n \bar{P}_0^{[\gamma_n]}(\cdot | x^m, \tau = n),
\end{aligned}$$

where, for the case $m = 0$, (*) follows from (3), which can be verified to be still valid in our generalized setting. The case $m > 0$ follows in exactly the same way, by shifting the data by m places (so that the new x_1 becomes what was x_{m+1} , and treating, for $k = 0, 1$, $\pi(H_k | x^m)$ as the priors for this shifted data problem, and then applying the above with $m = 0$).

We have shown that the Radon-Nikodym derivative $\frac{d\bar{P}_1^{[\gamma_n]}(\cdot | x^m)}{d\bar{P}_0^{[\gamma_n]}(\cdot | x^m)}$ at γ_n is given by $\gamma_n \cdot \frac{\bar{P}_0(\tau=n|x^m)\pi(H_1|x^m)}{\bar{P}_1(\tau=n|x^m)\pi(H_0|x^m)}$, which is what we had to show. \square

Proof. [of Lemma 7] Let A' denote the event $V_n = 1$ and Let $A \subset \mathbb{R}_{>0}$ be a Borel measurable set. We can write β_n as a function $\beta_n(U_n)$. With this notation, we have:

$$\begin{aligned}
P_{1,g}^{[\beta_n]}(A | A') &= \int_{\mathbb{R}_{>0}} \mathbb{1}_{\{A\}} dP_{1,g}^{[\beta_n]}(\cdot | A') \\
&= \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} dP_{1,g}^{[U_n]}(\cdot | A') \tag{29} \\
&= \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} \frac{dP_{1,g}^{[U_n]}(\cdot | A')}{dP_{0,g}^{[U_n]}(\cdot | A')} dP_{1,g}^{[U_n]}(\cdot | A') \\
&= \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} \mathbb{1}_{\{A'\}} \frac{dP_{1,g}^{[U_n]}(\cdot | A')}{dP_{0,g}^{[U_n]}(\cdot | A')} dP_{1,g}^{[U_n]}(\cdot | A') \\
&= \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} \mathbb{1}_{\{A'\}} \frac{P_{0,g}^{(n)}(A')}{P_{1,g}^{(n)}(A')} \frac{dP_{1,g}^{[U_n]}}{dP_{0,g}^{[U_n]}} dP_{0,g}^{[U_n]}(\cdot | A') \\
&= \frac{P_{0,g}^{(n)}(A')}{P_{1,g}^{(n)}(A')} \cdot \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} \mathbb{1}_{\{A'\}} \beta_n(U_n) dP_{0,g}^{[U_n]}(\cdot | A') \tag{30} \\
&= \frac{P_{0,g}^{(n)}(A')}{P_{1,g}^{(n)}(A')} \cdot \int_{\mathcal{U}_n} \mathbb{1}_{\{\beta_n(U_n) \in A\}} \beta_n(U_n) dP_{0,g}^{[U_n]}(\cdot | A') \\
&= \frac{P_{0,g}^{(n)}(A')}{P_{1,g}^{(n)}(A')} \cdot \int_{\mathbb{R}_{>0}} \mathbb{1}_{\{A\}} t dP_{0,g}^{[\beta_n]} A'(t),
\end{aligned}$$

where step (29) holds because β_n is \mathcal{G}_n -measurable and ((30)) follows from ((25)). We have shown that $\frac{P_{0,g}^{(n)}(A')}{P_{1,g}^{(n)}(A')} \cdot t$ is equal to the Radon-Nikodym derivative $\frac{dP_{1,g}^{[\beta_n]}(\cdot | V_n=1)}{dP_{0,g}^{[\beta_n]}(\cdot | V_n=1)}$, which is what we had to prove. \square