

Optimal Stopping for Interval Estimation in Bernoulli Trials

Tony Yaacoub, George V. Moustakides, *Senior Member, IEEE* and Yajun Mei

Abstract—We propose an optimal sequential methodology for obtaining confidence intervals for a binomial proportion θ . Assuming that an i.i.d. sequence of Bernoulli(θ) trials is observed sequentially, we are interested in designing (i) a stopping time T that will decide the best time to stop sampling the process, and (ii) an optimum estimator $\hat{\theta}_T$ that will provide the optimum center of the interval estimate of θ . We follow a semi-Bayesian approach, where we assume that there exists a prior distribution for θ , and our goal is to minimize the average number of samples while we guarantee a minimal specified coverage probability level. The solution is obtained by applying standard optimal stopping theory and computing the optimum pair $(T, \hat{\theta}_T)$ numerically. Regarding the optimum stopping time component T , we demonstrate that it enjoys certain very interesting characteristics not commonly encountered in solutions of other classical optimal stopping problems. In particular, we prove that, for a particular prior (Beta density), the optimum stopping time is always bounded from above and below; it needs to first accumulate a sufficient amount of information before deciding whether or not to stop, and it will always terminate before some finite deterministic time. We also conjecture that these properties are present with any prior. Finally, we compare our method with the optimum fixed-sample-size procedure as well as with existing alternative sequential schemes.

Index Terms—Sequential estimation, confidence intervals, binomial proportion, optimal stopping, sequential analysis.

I. INTRODUCTION

INTERVAL estimation of a binomial proportion θ is one of the most basic problems in statistics, with many important real-world applications. Some classical applications include interval estimation of the prevalence of a rare disease [1]; interval estimation of the overall response rate in clinical trials [2]; and accuracy assessment in remote sensing [3]. In these applications, the sample size is fixed in advance, and a *confidence interval* for θ is obtained. There is an extensive

Manuscript received November 18, 2017; revised August 3, 2018; accepted November 24, 2018. This work was supported by the US National Science Foundation under Grant CIF 1513373 through Rutgers University and under Grant CMMI 1362876 and DMS 1830344 through Georgia Institute of Technology.

T. Yaacoub was with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA. E-mail: tryaacoub@gmail.com.

G. V. Moustakides is with the Department of Computer Science, Rutgers University, Piscataway, NJ, USA and the Electrical and Computer Engineering Department, University of Patras, Rion, Greece. E-mail: george.moustakides@rutgers.edu, moustaki@upatras.gr. Website: <https://www.cs.rutgers.edu/~gm463/>.

Y. Mei is with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA. E-mail: ymei@isye.gatech.edu. Website: <http://www.isye.gatech.edu/~ymei/>.

Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

bibliography regarding derivations of confidence intervals for θ when the sample size is fixed. Perhaps the most widely known in this category is Wald's interval, which takes the form $\hat{\theta}_T \pm z_{\alpha/2} \sqrt{\hat{\theta}_T(1 - \hat{\theta}_T)/T}$, where T is the fixed sample size, $1 - \alpha$ expresses the desired coverage probability, $\hat{\theta}_T$ is the sample mean of θ and $z_{\alpha/2}$ satisfies $Q(z_{\alpha/2}) = \alpha/2$ with $Q(x)$ denoting the complementary cumulative density function (cdf) of a standard $N(0, 1)$ Gaussian random variable. This confidence interval is derived based on the asymptotic normality of $\hat{\theta}_T$ and, therefore, exhibits poor behavior when $T\theta(1 - \theta)$ is small [4]–[7]. Several efforts to improve Wald's classical method are reported in [4], [8]–[12]. There are also Bayesian-based techniques [5], [13], [14] while [4]–[7], [15] give interesting surveys that evaluate the relative performance of the above methods. Finally we mention that [16] provides explicit formulas for the required sample size that can guarantee a prescribed coverage probability for the Clopper-Pearson [9] method.

In many modern applications, sampling observations is costly and time consuming. Therefore, there is a desire to limit the sampling size without compromising the quality of the interval estimate. For instance, in automatic fraud detection in finance, one needs to manually go through the “suspect” financial transactions that are automatically detected as fraudulent by some machine learning or other computer algorithm. Since the manual process is expensive in terms of labor and cost, it is desirable to quickly estimate, with high confidence, what percentage of the suspect transactions are truly fraudulent. A different motivating application is in Statistical Model Checking, where with an approximate verification method, one overcomes the state space explosion problem for probabilistic systems by the use of Monte Carlo simulations. Given an executable stochastic system, we seek to verify that the system satisfies a particular property via simulation; and we desire to estimate the probability θ by which the system actually satisfies the property in question. The goal is to estimate θ within acceptable margins of error and confidence (see [17] and references therein). Because Monte Carlo simulations very often tend to require extensive time and computing power, it is advantageous to reduce their number assuring, at the same time, satisfactory quality levels for the corresponding estimate. The sequential version of the interval estimation aims exactly at reducing the sample size by selecting it to be random and, in particular, a stopping time controlled by the observations themselves. The literature focusing on the sequential setup of the problem is limited compared to its fixed sample-size counterpart (see [18]–[20]).

However, none of these articles that focus on sequential fixed-width confidence intervals for θ is able to claim optimality of their schemes.

The objective of our current work is to offer optimum *sequential* methods for interval estimation of θ , with the quality of the estimate expressed through the *coverage probability*. In addition to deriving the optimum scheme, we will also demonstrate some very uncommon but highly interesting properties of the optimum solution. These properties are not encountered in optimum sequential schemes derived for other well known sequential problems (i.e. sequential hypothesis testing). In particular, we show that the optimum stopping time is bounded from above for a rich class of priors, proving that the stopping time will always terminate in finite time. Furthermore, we show that the optimum stopping time is bounded from below for the Beta density prior, which implies that we first need a sufficient number of samples before we start applying our stopping rule. We also develop an efficient recursion to obtain numerically the coverage probability and the expected number of samples for the optimum scheme without using any Monte Carlo simulations.

We must also add that our methodology for interval estimation exhibits similarities with the work developed in [21] for point estimation. The main similarity is using Bayesian risk minimization and reducing this minimization to a standard optimal stopping problem as we explain in Section II. The second similarity is showing that the optimum stopping time is bounded from above. However, the focus in [21] is on the actual point estimate of θ with the adopted risk function being a variation of the classical mean square error; as a result, the corresponding optimum solution has a simple, explicit form. Note that there is a significant difference between point and interval estimation. In point estimation, the objective is to estimate an unknown population parameter with a single value based on the observed data. On the other hand, in interval estimation, the objective is to obtain a range of values within which, we believe, the true population parameter lies with high coverage probability. In our work, as we pointed out, we focus on confidence intervals and coverage probabilities. In this case, the width of the interval and the confidence coefficient jointly provide a sense of precision and accuracy. We unify these criteria under an appropriately defined indicator loss function that leads to the risk function being the complementary coverage probability. As it turns out, this difference makes our derivations and proofs far more complicated, requiring original analytical methodology. This becomes particularly apparent when our optimum solution does not have an explicit analytic form, and we attempt to establish the validity of the unique properties, mentioned before, that characterize our optimum solution.

The remainder of this article is organized as follows. In Section II we discuss our proposed framework for interval estimation for θ and propose a well-defined optimization problem and discuss its general solution. In Section III we focus on the computational aspects of the optimum scheme and the unique properties that they characterize it. In Section IV we compare the proposed scheme against the fixed-sample-size and two existing sequential methods in the literature. Finally,

Section V contains our conclusions.

II. PROPOSED FRAMEWORK

We observe sequentially an independent and identically distributed (i.i.d.) process X_1, X_2, \dots of Bernoulli random variables with $X_t \in \{0, 1\}$ and $P(X_t = 1) = \theta = 1 - P(X_t = 0)$ for some unknown $\theta \in [0, 1]$. The goal is to provide a confidence interval for θ . We are interested in confidence intervals of fixed width equal to $2h$ for some pre-specified $h \in (0, 1/2)$. We would also like our scheme to be able to guarantee a coverage probability equal to $1 - \alpha$, where $\alpha \in (0, 1)$ is given. Our scheme consists of a pair $(T, \hat{\theta}_T)$, that is, a stopping time T and a *mid-point estimator*¹ $\hat{\theta}_T$, where T is adapted to the observation history (filtration generated by the observations) and $\hat{\theta}_T$ is a function of the observations accumulated up to the time of stopping T . We would like to solve the following constrained optimization problem for the optimum pair

$$\inf_{T, \hat{\theta}_T} E[T|\theta], \text{ subject to: } P(|\hat{\theta}_T - \theta| > h|\theta) \leq \alpha, \quad (1)$$

where the desired interval estimate is $[\hat{\theta}_T - h, \hat{\theta}_T + h]$ (with the two ends cropped at 0 and 1, respectively, whenever they exceed the two limits) and where $P(\cdot|\theta)$ and $E[\cdot|\theta]$ denote probability and expectation for *given* θ .

Although (1) seems as the ideal formulation, it unfortunately targets an infeasible goal. We note that we are asking for the pair $(T, \hat{\theta}_T)$ to minimize the average number of samples for *every value* of the parameter θ . In other words, it is impossible to find a scheme that is optimal for each and every θ . In order to be able to find a solution that has a well-defined form of optimality, we adopt a semi-Bayesian approach² and assume that a prior $\pi(\theta)$ for θ is available. This allows for the following modification of the previous constrained optimization

$$\inf_{T, \hat{\theta}_T} E[T], \text{ subject to: } P(|\hat{\theta}_T - \theta| > h) \leq \alpha \quad (2)$$

where $P(\cdot)$ and $E[\cdot]$ denote probability and expectation including *averaging over* θ with the help of the prior; $P(\cdot) = \int_0^1 P(\cdot|\theta)\pi(\theta)d\theta$ and $E[\cdot] = \int_0^1 E[\cdot|\theta]\pi(\theta)d\theta$, respectively.

Remark 1. We must emphasize that the constraint in (2) does not guarantee that the desired coverage probability will also hold for each individual θ , namely $P(|\hat{\theta}_T - \theta| > h|\theta) \leq \alpha$, a property which is particularly desirable in practice. Perhaps, a more meaningful problem to consider in place of (2) would have been

$$\inf_{T, \hat{\theta}_T} \sup_{\theta} E[T|\theta], \text{ subject to: } \sup_{\theta} P(|\hat{\theta}_T - \theta| > h|\theta) \leq \alpha, \quad (3)$$

that assures a coverage probability of at least $1 - \alpha$ for *every* θ . In principle, this setup has an optimal solution and does not

¹The estimate $\hat{\theta}_T$ does not have the meaning of a classical parameter estimator. It is the mid-point of the confidence interval $[\hat{\theta}_T - h, \hat{\theta}_T + h]$ and does not necessarily constitute an efficient estimate of θ .

²The term “semi-Bayesian” is used because our setup involves two different components where one is optimized while the other is constrained, unlike full-Bayesian approaches that combine all terms into a single performance measure.

require a prior for θ , since it is a min-max problem. However, deriving the optimal solution for this min-max formulation is usually very complicated, and such solution has not yet been derived even for the fixed-sample-size case. Instead, we focus on (2) as the optimum scheme we would like to develop, but in our numerical examples, we will also evaluate $E[T|\theta]$ in terms of $\sup_{\theta} P(|\hat{\theta}_T - \theta| > h|\theta) \leq \alpha$.

Let $c > 0$ denote a Lagrange multiplier that we use to combine the two terms in (2) into a single cost function $J(T, \hat{\theta}_T) = cE[T] + P(|\hat{\theta}_T - \theta| > h)$, and consider the *unconstrained* optimization problem

$$\inf_{T, \hat{\theta}_T} J(T, \hat{\theta}_T) = \inf_{T, \hat{\theta}_T} \left\{ cE[T] + P(|\hat{\theta}_T - \theta| > h) \right\}. \quad (4)$$

Note that for any estimator $\hat{\theta}$ of θ , the complementary coverage probability $P(|\hat{\theta} - \theta| > h|\theta)$ can be written as a risk function in the Bayesian setting if we define the loss function to be $L(\hat{\theta}, \theta) = \mathbb{1}_{\{|\hat{\theta} - \theta| > h\}}$, where $P(|\hat{\theta} - \theta| > h|\theta) = E[L(\hat{\theta}, \theta)|\theta]$. This allows us to combine the frequentist idea of confidence intervals into a Bayesian setting; the constraint in (1) becomes $E[L(\hat{\theta}_T, \theta)|\theta] \leq \alpha$, and the constraint in (2) becomes $E[L(\hat{\theta}_T, \theta)] = \int E[L(\hat{\theta}_T, \theta)|\theta] \pi(\theta) d\theta \leq \alpha$. By doing so, (4) can be solved using standard optimal stopping theory. We will first identify the solution to (4) and then demonstrate that a proper selection of c can also solve the constrained problem in (2).

A. The Unconstrained Problem

We start by considering the classical Bayes estimation problem for fixed sample size t

$$\inf_{\hat{\theta}_t} P(|\hat{\theta}_t - \theta| > h). \quad (5)$$

If we observe $\mathcal{F}_t = \sigma\{X_1, \dots, X_t\}$ then, given that $\{X_t\}$ is i.i.d. Bernoulli(θ), the probability to obtain a specific combination of samples given θ is equal to $\theta^{S_t} (1 - \theta)^{t - S_t}$, where $S_t = \sum_{k=1}^t X_k$ is the number of “successes” up to time t . This implies that the posterior probability density of θ given the observations can be written as

$$\pi_t(\theta|\mathcal{F}_t) = \pi_t(\theta|S_t) = \frac{\theta^{S_t} (1 - \theta)^{t - S_t} \pi(\theta)}{\int_0^1 \theta^{S_t} (1 - \theta)^{t - S_t} \pi(\theta) d\theta}. \quad (6)$$

From Bayesian estimation theory [22, Page 142], we have that the optimization in (5) is achieved by the following Bayes estimator

$$\begin{aligned} \hat{\theta}_t(S_t) &= \arg \inf_{\hat{\theta}_t} P(|\hat{\theta}_t - \theta| > h|\mathcal{F}_t) \\ &= \arg \sup_{\hat{\theta}_t} \int_{\max\{\hat{\theta}_t - h, 0\}}^{\min\{\hat{\theta}_t + h, 1\}} \pi_t(\theta|S_t) d\theta, \end{aligned} \quad (7)$$

yielding the corresponding optimum conditional complementary coverage probability

$$\begin{aligned} \mathcal{C}_t(S_t) &= \inf_{\hat{\theta}_t} P(|\hat{\theta}_t - \theta| > h|\mathcal{F}_t) \\ &= 1 - \sup_{\hat{\theta}_t} \int_{\max\{\hat{\theta}_t - h, 0\}}^{\min\{\hat{\theta}_t + h, 1\}} \pi_t(\theta|S_t) d\theta \\ &= 1 - \int_{\max\{\hat{\theta}_t(S_t) - h, 0\}}^{\min\{\hat{\theta}_t(S_t) + h, 1\}} \pi_t(\theta|S_t) d\theta. \end{aligned} \quad (8)$$

From (7) and (8) we observe that both quantities $\hat{\theta}_t(S_t), \mathcal{C}_t(S_t)$ are \mathcal{F}_t -measurable and, more precisely, functions of S_t . For known prior $\pi(\theta)$, we can, at least numerically, compute the Bayes estimate and the corresponding optimum conditional complementary coverage probability for each combination of integer pair (t, S_t) .

Remark 2. By focusing on (7), we can make a small but interesting observation: Regarding the Bayes estimate $\hat{\theta}_t(S_t)$ it is easy to verify that

$$h \leq \hat{\theta}_t(S_t) \leq 1 - h. \quad (9)$$

Indeed, this is clear, because if in (7) we select $\hat{\theta}_t < h$ or $\hat{\theta}_t > 1 - h$, this will yield an inferior cost compared to the selection $\hat{\theta}_t = h$ or $\hat{\theta}_t = 1 - h$, respectively. The implication of this observation is that $\hat{\theta}_t(S_t)$ will be *biased and inconsistent* when considered as an estimate of the true parameter θ , at least for values of θ outside the interval $[h, 1 - h]$. As we mentioned, the correct meaning of this quantity is that it constitutes the *mid-point* of the confidence interval $[\hat{\theta}_t(S_t) - h, \hat{\theta}_t(S_t) + h]$ with the latter enjoying, for each fixed t , the largest possible coverage probability.

Consider now the optimization in (4) which will be performed in two steps: First we fix the stopping time T and minimize $J(T, \hat{\theta}_T)$ with respect to $\hat{\theta}_T$; the resulting expression is then minimized, during the second step, over T in order to obtain the optimum pair. We have the following lemma that addresses the first problem.

Lemma 1. *For any given stopping time T , we have*

$$\begin{aligned} J(T, \hat{\theta}_T) &= cE[T] + P(|\hat{\theta}_T - \theta| > h) \\ &\geq E[cT + \mathcal{C}_T] = J(T), \end{aligned} \quad (10)$$

with equality when we apply the corresponding Bayesian estimator $\hat{\theta}_T = \hat{\theta}_T$ at the time of stopping.

Proof. The proof depends on whether $E[T] = \infty$ or $E[T] < \infty$. The details are presented in the Appendix. \square

A side-product of Lemma 1, as it can be verified from the corresponding proof in the Appendix, is the fact that the Bayesian estimator is not only optimum for fixed sample size, but it retains its optimality property when the sample size is controlled by any stopping time T adapted to the observations.

Using (10) from Lemma 1, we are now left with the optimization of the stopping time T . Assuming that N is an integer which is sufficiently large, we consider the following optimization over stopping times that are bounded by N

$$\inf_{0 \leq T \leq N} J(T) = \inf_{0 \leq T \leq N} E[cT + \mathcal{C}_T]. \quad (11)$$

This is a classical *finite horizon* optimal stopping problem with cost per sample equal to c and cost for stopping at t equal to \mathcal{C}_t . Of course, it is only natural to wonder why we limited our analysis to finite horizons instead of considering the more classical *infinite horizon* version. As we will see in the sequel, for the most common prior we will be able to demonstrate that the infinite horizon assumption is completely unnecessary. Indeed, the optimum stopping time will turn out to be bounded by a deterministic quantity, suggesting that by limiting ourselves to a (sufficiently large) finite horizon, we do not suffer any performance loss.

In order to solve the optimization problem defined in (11), we follow the classical optimal stopping theory [23]. For $t = 0, 1, \dots, N$ define the sequence of optimal *average residual costs*

$$\mathcal{V}_t = \inf_{t \leq T \leq N} \mathbb{E}[c(T - t) + \mathcal{C}_T | \mathcal{F}_t], \quad (12)$$

then we have

$$\mathcal{V}_t = \min\{\mathcal{C}_t, c + \mathbb{E}[\mathcal{V}_{t+1} | \mathcal{F}_t]\}, \quad t = N, \dots, 1, 0, \quad (13)$$

with the backward recursion initialized with $\mathcal{V}_{N+1} = 1$. Regarding this last selection, it produces $\mathcal{V}_N = \mathcal{C}_N$ since the latter is a probability. In fact, this is exactly what the optimum residual cost at N must be, because if we have not stopped before N , then we necessarily stop at N and this produces cost \mathcal{C}_N (simply the cost of stopping at N). The total optimum cost is expressed through \mathcal{V}_0 , namely $\mathcal{V}_0 = \inf_{0 \leq T \leq N} \mathcal{J}(T)$. The next lemma specifies in more detail the recursion in (13).

Lemma 2. *Consider the recursion in (13) then, the optimal residual cost $\mathcal{V}_t, t = N, \dots, 0$ is a function $\mathcal{V}_t(S_t)$ of S_t and therefore \mathcal{F}_t -measurable. Furthermore, (13) can be written as*

$$\mathcal{V}_t(S_t) = \min\{\mathcal{C}_t(S_t), c + \tilde{\mathcal{V}}_t(S_t)\}, \quad t = N, \dots, 0, \quad (14)$$

where $\tilde{\mathcal{V}}_t(S_t)$ expresses the optimum average residual cost to continue, satisfying

$$\tilde{\mathcal{V}}_t(S_t) = g_{t+1}(S_t) \mathcal{V}_{t+1}(S_t + 1) + (1 - g_{t+1}(S_t)) \mathcal{V}_{t+1}(S_t), \quad (15)$$

$$g_{t+1}(S_t) = \mathbb{P}(X_{t+1} = 1 | \mathcal{F}_t) = \frac{\int_0^1 \theta^{S_t+1} (1 - \theta)^{t-S_t} \pi(\theta) d\theta}{\int_0^1 \theta^{S_t} (1 - \theta)^{t-S_t} \pi(\theta) d\theta}. \quad (16)$$

Finally, if the prior $\pi(\theta)$ is symmetric around $\frac{1}{2}$ then the functions $\mathcal{C}_t(S_t), \mathcal{V}_t(S_t), \tilde{\mathcal{V}}_t(S_t)$ are symmetric with respect to S_t around the value $\frac{t}{2}$.

Proof. The validity of this lemma is straightforward and can be easily established using induction. We therefore give no further details. \square

Once the sequence of optimal residual costs has been obtained through the solution of (14), it is then immediate to define the optimum stopping time T_o that solves the minimization problem in (11). Again, optimal stopping theory [23] suggests that

$$\begin{aligned} T_o &= \inf\{0 \leq t \leq N : \mathcal{V}_t(S_t) = \mathcal{C}_t(S_t)\} \\ &= \inf\{0 \leq t \leq N : \mathcal{C}_t(S_t) \leq c + \tilde{\mathcal{V}}_t(S_t)\}. \end{aligned} \quad (17)$$

In other words, when the optimum residual cost $\mathcal{V}_t(S_t)$ matches, for the first time, the cost for stopping $\mathcal{C}_t(S_t)$ or, equivalently, the cost of stopping is smaller than the residual cost of continuing, this is when we stop. Since the functions involved depend on S_t , this quantity can serve as our test statistic and we can express the stopping rule in (17) in terms of S_t . Specifically, for each time t , we can find the sampling region $\Omega_t = \{0 \leq S_t \leq t : \mathcal{V}_t(S_t) < \mathcal{C}_t(S_t)\} = \{0 \leq S_t \leq t : c + \tilde{\mathcal{V}}_t(S_t) < \mathcal{C}_t(S_t)\}$ with $\Omega_N = \emptyset$, and we can equivalently define the stopping time as $T_o = \inf\{0 \leq t \leq N : S_t \notin \Omega_t\}$.

B. The Constrained Problem

Let us now turn to the constrained problem in (2) which we can solve with the results we have so far. We will show that (2) can be recovered as an instance of the unconstrained version (4) corresponding to a special selection of the Lagrange multiplier c . Our result is summarized in the following theorem.

Theorem 1. *For the solution of (2) we distinguish two cases, depending on the value of $\mathcal{C}_0 = \mathbb{P}(|\hat{\theta}_0 - \theta| > h)$, where $\hat{\theta}_0 = \arg \inf_{\hat{\theta}_0} \mathbb{P}(|\hat{\theta}_0 - \theta| > h)$, i.e., the Bayes estimator in (7) at $t = 0$ with $\pi_0(\theta | S_0) = \pi(\theta)$.*

- i) *If $\mathcal{C}_0 \leq \alpha$, then the optimum is to stop without taking any samples, i.e. $T_o = 0$ and use as mid-point of the optimum confidence interval the value $\hat{\theta}_0$ that is based only on the prior $\pi(\theta)$.*
- ii) *If $\mathcal{C}_0 > \alpha$, then for any horizon $N \geq N_\alpha$ where N_α satisfies $\mathbb{P}(|\hat{\theta}_{N_\alpha} - \theta| > h) < \alpha$, there exists Lagrange multiplier $c_* \in (0, 1)$, independent from N , with c_* being the solution to $\mathbb{P}(|\hat{\theta}_{T_o(c_*)} - \theta| > h) = \alpha$, such that the solution of (4) is also the solution to (2).*

Proof. The proof of this theorem is presented in the Appendix. \square

III. PROPERTIES OF THE OPTIMUM SOLUTION

If we fix the value N of the horizon and the cost per sample c , we can then compute the mid-points $\{\{\hat{\theta}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ of the confidence intervals from (7). Assuming that $\pi(\theta)$ is continuous, candidates for $\hat{\theta}_t(S_t)$ can be obtained from the solution of the following equation which we obtain by differentiating (7) with respect to $\hat{\theta}_t$

$$\begin{aligned} &(\hat{\theta}_t + h)^{S_t} (1 - \hat{\theta}_t - h)^{t-S_t} \pi(\hat{\theta}_t + h) \\ &- (\hat{\theta}_t - h)^{S_t} (1 - \hat{\theta}_t + h)^{t-S_t} \pi(\hat{\theta}_t - h) = 0. \end{aligned} \quad (18)$$

The previous equation has clearly a solution in the interval $[h, 1 - h]$ when $0 < S_t < t$ with the corresponding value providing a (local) extremum for the coverage probability. To these candidate mid-points we must include the two end points $h, 1 - h$ since the global maximum can occur at the two ends as well. Therefore, we need to examine which of these cases provides the best coverage probability and select the corresponding value as our optimum mid-point $\hat{\theta}_t(S_t)$. When $S_t = 0, t$ it is possible (18) not to have any solution in $[h, 1 - h]$. In this case, $\hat{\theta}_t(0)$ and $\hat{\theta}_t(t)$ are equal to one of the two end values h or $1 - h$. Having identified the optimum

mid-points $\{\{\hat{\theta}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$, we apply (8) to compute the corresponding optimum complementary conditional coverage probabilities $\{\{\mathcal{C}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$.

The next step consists in computing $\{\{g_{t+1}(S_t)\}_{S_t=0}^t\}_{t=0}^N$ for $t = 0, \dots, N$ and $S_t = 0, \dots, t$ with numerical integration. Once we have available $\{\{\mathcal{C}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ and $\{\{g_{t+1}(S_t)\}_{S_t=0}^t\}_{t=0}^N$, we can then use them in the backward recursion (14) to find the sequence $\{\{\tilde{\mathcal{V}}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ and the optimum residual cost sequence $\{\{\mathcal{V}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$. To identify the stopping rule, according to (17) we must compare the two sequences $\{\{\mathcal{C}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$, $\{\{\mathcal{V}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ element-by-element. At coordinates (t, S_t) where the sequences differ, we decide to continue sampling; whereas if they are equal, we decide to stop. This generates the sequence of sampling regions $\{\Omega_t\}_{t=0}^N$. Equivalently, we can compare $\{\{\mathcal{C}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ with $\{\{c + \tilde{\mathcal{V}}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$, and whenever the first is no larger than the second, we stop, while we continue sampling in the opposite case.

We now present a conjecture that contains two significant claims for the optimum stopping time for the problem in (4) which we believe are valid for *any* prior $\pi(\theta)$. We were able to provide a proof for the first claim (Lemma 3) for a rich class of priors, and prove both claims (Theorem 2) providing also quantitative information when the prior is the Beta density. Regarding the latter case we should note that the Beta density is among the most popular priors for the problem we are considering in this work.

Conjecture. *For any prior $\pi(\theta)$ and sufficiently large horizon N the optimum stopping time T_o of the unconstrained problem in (4) enjoys the following two properties:*

- i). *There exists constant t_{up} depending only on c and not on N such that $T_o \leq t_{\text{up}}$.*
- ii). *For sufficiently small c there exists constant $t_{\text{lo}} \geq 1$ depending only on c and not on N such that $t_{\text{lo}} \leq T_o$.*

Below we present a general proof of property i) of the Conjecture under the following additional assumption: Define the maximal conditional variance

$$\begin{aligned} \sigma_t^2 &= \max_{0 \leq S_t \leq t} \mathbb{E}[(\theta - \mathbb{E}[\theta|S_t])^2|S_t] \\ &= \max_{0 \leq S_t \leq t} \int_0^1 (\theta - \mathbb{E}[\theta|S_t])^2 \pi_t(\theta|S_t) d\theta, \end{aligned} \quad (19)$$

where $\pi_t(\theta|S_t)$ is the posterior pdf defined in (6) and assume that $\sigma_t \rightarrow 0$ as $t \rightarrow \infty$. This forces the conditional variance to converge to 0 *uniformly* in S_t . It also implies that the posterior distribution $\pi_t(\theta|S_t)$ converges, uniformly, to a degenerate measure at a single point (often the true θ) as $t \rightarrow \infty$. This is clearly related to the *consistency* concept of posterior distributions in Bayesian statistics and is often considered a valid assumption (see [24]).

Lemma 3. *Let σ_t be defined as in (19) with $\lim_{t \rightarrow \infty} \sigma_t = 0$. Then for sufficiently large horizon there exists constant t_{up} depending only on c such that $T_o \leq t_{\text{up}}$, i.e. property i) in the Conjecture is true.*

Proof. The proof is a simple application of the Chebyshev inequality in combination with (19). Indeed we observe that

$$\begin{aligned} \mathcal{C}_t(S_t) &= \inf_{\hat{\theta}_t} \mathbb{P}(|\theta - \hat{\theta}_t| > h|\mathcal{F}_t) \leq \mathbb{P}(|\theta - \mathbb{E}[\theta|\mathcal{F}_t]| > h|\mathcal{F}_t) \\ &\leq \frac{1}{h^2} \mathbb{E}[(\theta - \mathbb{E}[\theta|S_t])^2|S_t] \leq \frac{\sigma_t^2}{h^2}. \end{aligned} \quad (20)$$

Since $\sigma_t \rightarrow 0$ as $t \rightarrow \infty$, there exists N such that $\mathcal{C}_N \leq \frac{\sigma_N^2}{h^2} \leq c$ and, therefore, from (14) we conclude that $\mathcal{C}_N \leq c + \tilde{\mathcal{V}}_N$, which suggests that we will necessarily stop at N for any value of S_N . Quantity t_{up} is the smallest N for which this is true. \square

Remark 3. The assumption $\lim_{t \rightarrow \infty} \sigma_t = 0$ does not hold for all prior distribution. A counterexample where it fails is when the prior is a two-point probability mass function, say $\mathbb{P}(\theta = 0.4) = \mathbb{P}(\theta = 0.6) = 0.5$. However, even for this case the Conjecture might still be valid since the requirement $\mathcal{C}_N(S_N) \leq \sigma_N^2/h^2 < c$ used in our proof, is only sufficient for the validity of our claim.

An interesting example where the assumption holds is when the prior is the Beta density $\pi(\theta) = \text{Beta}(\theta, p, q)$, where

$$\text{Beta}(\theta, p, q) = \frac{\theta^{p-1}(1-\theta)^{q-1}}{\int_0^1 \theta^{p-1}(1-\theta)^{q-1} d\theta}, \quad p, q > 0. \quad (21)$$

To see this, we note that the posterior pdf is of the same type, namely $\pi(\theta|S_t) = \text{Beta}(\theta, p + S_t, t - S_t + q)$, and thus the maximal conditional variance in (19) becomes

$$\begin{aligned} \sigma_t^2 &= \max_{0 \leq S_t \leq t} \frac{(p + S_t)(t - S_t + q)}{(t + p + q)^2(t + p + q + 1)} \\ &\leq \frac{1}{4(t + p + q + 1)}, \end{aligned} \quad (22)$$

where the equality is attainable when $S_t = (t + q - p)/2$ is an integer. Clearly, for fixed $p, q > 0$ we have $\sigma_t \rightarrow 0$ as $t \rightarrow \infty$, and thus the assumption of Lemma 3 holds. Moreover, by the proof of Lemma 3, the optimum stopping time satisfies $T_o \leq \max\{0, 1/(4h^2c) - p - q - 1\}$ for all $c > 0$. This bound is of the order of c^{-1} . In Theorem 2, Section III-B, by applying a more advanced analysis, we will be able to improve it and provide an alternative estimate which is of the order of $|\log(c)|$ for the case of the symmetric prior $p = q$.

Remark 4. Property i) of the Conjecture suggests that the number of samples, under the optimum scheme, will never exceed the value t_{up} even if we allow the horizon to grow without limit. This interesting and uncommon characteristic was also observed in [21] but with cost function a variance of the classical mean square error. However, what is more intriguing in our conjecture is property ii), namely that we need first to accumulate a sufficient volume of information before we start asking ourselves whether we should stop sampling or not. This is an extremely uncommon feature and, to our knowledge, has never been reported before in Sequential Analysis as a property of optimum schemes. As we claim in our conjecture, we believe that both properties are valid for any prior $\pi(\theta)$. Fortunately, as we mentioned before, this double

claim is not without solid evidence. Indeed with Theorem 2, we demonstrate its validity when the prior is the symmetric Beta density.

A. Performance Evaluation

What we presented so far allows for the determination of the stopping rule of the proposed scheme. We would like now to compute its performance but also the performance of any stopping time which uses S_t as its test statistic and is defined in terms of a sequence of sampling regions $\{\Omega_t\}$ in terms of $\{S_t\}$. In particular, we are interested in computing $E[T|\theta]$, $E[T]$, $P(|\hat{\theta}_T - \theta| \leq h|\theta)$ and $P(|\hat{\theta}_T - \theta| \leq h)$. Of course, we could obtain these quantities using Monte-Carlo simulations, but it is also possible to determine them numerically. The following lemma provides the necessary formulas.

Lemma 4. *Let the stopping time T be bounded by N having as test statistic the process $\{S_t\}$. Assume for each t that Ω_t denotes the sampling region. Suppose also that for the combination (t, S_t) the scheme provides the mid-point estimate $\hat{\theta}_t(S_t)$ and the corresponding conditional complementary coverage probability $C_t(S_t) = P(|\hat{\theta}_t(S_t) - \theta| > h|\mathcal{F}_t)$. For $t = N-1, \dots, 0$, we then define the following backward recursions that must be applied for $S_t = 0, 1, \dots, t$*

$$U_t(S_t) = 1 + \theta \mathbb{1}_{\{S_t+1 \in \Omega_{t+1}\}} U_{t+1}(S_t + 1) + (1 - \theta) \mathbb{1}_{\{S_t \in \Omega_{t+1}\}} U_{t+1}(S_t), \quad (23)$$

$$\bar{U}_t(S_t) = 1 + g_{t+1}(S_t) \mathbb{1}_{\{S_t+1 \in \Omega_{t+1}\}} \bar{U}_{t+1}(S_t + 1) + (1 - g_{t+1}(S_t)) \mathbb{1}_{\{S_t \in \Omega_{t+1}\}} \bar{U}_{t+1}(S_t), \quad (24)$$

$$W_t(S_t) = \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \mathbb{1}_{\{S_t \notin \Omega_t\}} + \{\theta W_{t+1}(S_t + 1) + (1 - \theta) W_{t+1}(S_t)\} \mathbb{1}_{\{S_t \in \Omega_t\}}, \quad (25)$$

$$\bar{W}_t(S_t) = C_t(S_t) \mathbb{1}_{\{S_t \notin \Omega_t\}} + \{g_{t+1}(S_t) \bar{W}_{t+1}(S_t + 1) + (1 - g_{t+1}(S_t)) \bar{W}_{t+1}(S_t)\} \mathbb{1}_{\{S_t \in \Omega_t\}}, \quad (26)$$

where $g_{t+1}(S_t)$ is defined in (16) and the four recursions are initialized with $U_N(S_N) = \bar{U}_N(S_N) = 0$, $W_N(S_N) = \mathbb{1}_{\{|\hat{\theta}_N - \theta| > h\}}$, $\bar{W}_N(S_N) = C_N(S_N)$, $\Omega_N = \emptyset$. Then, $E[T|\theta] = U_0(S_0)$, $E[T] = \bar{U}_0(S_0)$, $P(|\hat{\theta}_T - \theta| > h|\theta) = W_0(S_0)$ and $P(|\hat{\theta}_T - \theta| > h) = \bar{W}_0(S_0)$.

Proof. The validity of these expressions is established in the Appendix. \square

The applicability of Lemma 4 is clearly not limited to the proposed scheme but can be used to compute the performance of the fixed-sample-size and of other sequential alternatives that we intend to compare against the method we have developed.

B. Beta Density as Prior

Let us now find the particular form of our scheme when we adopt as our prior the Beta density $\pi(\theta) = \text{Beta}(\theta, a, a)$, where $\text{Beta}(\theta, p, q)$ is defined in (21).

We observe that the selection $a = 1$ in the prior corresponds to the uniform density in $[0, 1]$. It is now straightforward to verify that the posterior pdf accepts a similar form, namely

$$\pi(\theta|S_t) = \text{Beta}(\theta, a + S_t, a + t - S_t), \quad (27)$$

while the conditional complementary coverage probability at time t becomes

$$P(|\hat{\theta}_t - \theta| > h|\mathcal{F}_t) = 1 - I_{\min\{1, \hat{\theta}_t + h\}}(a + S_t, a + t - S_t) + I_{\max(0, \hat{\theta}_t - h)}(a + S_t, a + t - S_t), \quad (28)$$

where $I_x(p, q)$ is the incomplete Beta function (see [25, Page 944]) which is the cdf of $\text{Beta}(\theta, p, q)$.

The Bayes estimator, according to (18), can be found as the solution of the equation

$$\hat{\vartheta}_t = \arg \left\{ \hat{\theta}_t : \left(\frac{\hat{\theta}_t - h}{\hat{\theta}_t + h} \right)^{a + S_t - 1} = \left(\frac{1 - h - \hat{\theta}_t}{1 + h - \hat{\theta}_t} \right)^{a + t - S_t - 1} \right\}$$

corresponding to the root in the interval $[h, 1 - h]$. Such root always exists except when $a = 1$ and $S_t = 0$ or t . For these cases, $\hat{\vartheta}_t$ is equal to h or $1 - h$, depending on which value provides a larger conditional coverage probability. The resulting optimum conditional complementary coverage probability becomes

$$\mathcal{C}_t(S_t) = 1 - I_{\min\{1, \hat{\vartheta}_t + h\}}(a + S_t, a + t - S_t) + I_{\max(0, \hat{\vartheta}_t - h)}(a + S_t, a + t - S_t). \quad (29)$$

Finally, as indicated in (15) and (16), we need to find the probability $g_{t+1}(S_t)$, for which we have the following simple formula

$$g_{t+1}(S_t) = P(X_{t+1} = 1|\mathcal{F}_t) = \frac{\Gamma(S_t + a + 1)\Gamma(t + 2a)}{\Gamma(S_t + a)\Gamma(t + 2a + 1)} = \frac{S_t + a}{t + 2a}. \quad (30)$$

We can now compute the sequences $\{\{\mathcal{V}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$, $\{\{\tilde{\mathcal{V}}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ as explained in (14) and compare, element-by-element, $\{\{\mathcal{C}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ with $\{\{\mathcal{V}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ or $\{\{\tilde{\mathcal{C}}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ with $\{\{c + \tilde{\mathcal{V}}_t(S_t)\}_{S_t=0}^t\}_{t=0}^N$ to identify the sampling and stopping regions.

For the particular prior adopted in (21), as we mentioned before, the resulting optimum stopping time T_o enjoys the unique properties claimed in the Conjecture. The next theorem provides the necessary evidence.

Theorem 2. *The Conjecture is true when the prior is the Beta density $\pi(\theta) = \text{Beta}(\theta, a, a)$ with the optimum stopping time T_o satisfying $C_0 |\log(c)| \leq T_o \leq C_1 |\log(c)|$ for constants $C_0 < C_1$ that depend only on a and h .*

Proof. The proof is very technical and detailed in the Appendix. Unfortunately, the analytical techniques developed for the specific prior are not directly extendable to the general case. \square

Perhaps, it is worth mentioning the fact that from the proof of Theorem 2, we conclude that the two estimates for t_{up} in (45) and t_{lo} in (47) grow linearly in $|\log(c)|$ having

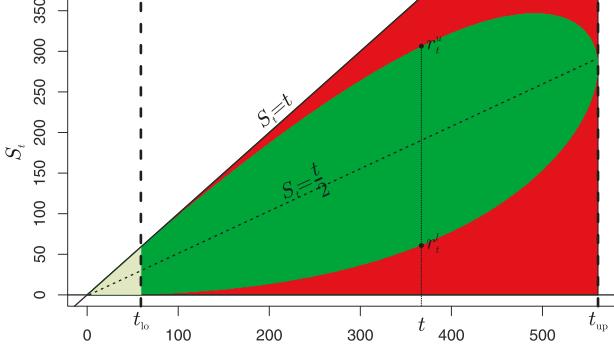


Fig. 1. Sampling (green) and stopping (red) regions for $a = 1$, $h = 0.05$ and $c = 0.0001$. Upper and lower bounds for optimum stopping time: $t_{lo} = 59$ and $t_{up} = 561$. No possibility of stopping (light green).

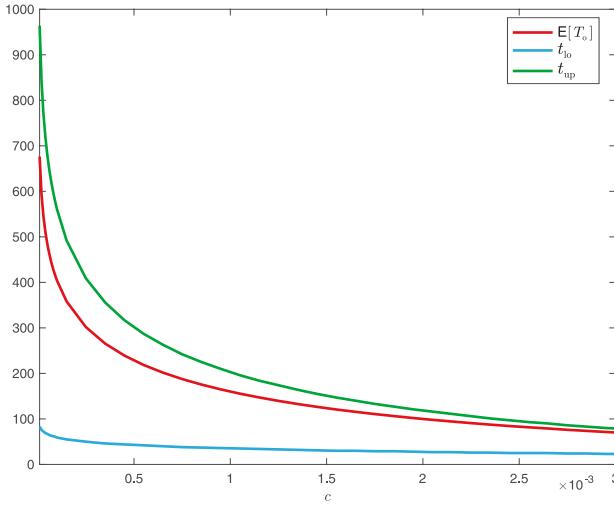


Fig. 2. Average sample size (red), lower t_{lo} (blue) and upper t_{up} limit (green), as functions of c for optimum stopping time T_o when $a = 1$ and $h = 0.05$.

drastically different multiplicative coefficients (C_0 of the order of $1/|\log(0.5 - h)|$ versus C_1 of the order of $1/(2h^2)$) and different offsets.

As an illustration for these properties we consider $a = 1$, $h = 0.05$, and $c = 0.0001$. Fig. 1 depicts the sampling (green) and the stopping (red) region in terms of the test statistic S_t . Both regions are clearly limited between the lines $S_t = t$ and $S_t = 0$. Even though we have marked a whole region in red, only the points that are next to the green region are actually accessible because S_t can increase at most by one unit as we go from t to $t + 1$. We can also see the two bounds $t_{up} = 561$ and $t_{lo} = 59$ for T_o . For $t \leq t_{lo}$ the light green region covers all points $0 \leq S_t \leq t$, thus identifying the time instances we can never stop. Also, we note that once we pass t_{up} we are in the stopping region suggesting that we must necessarily stop at t_{up} . For each $t_{lo} \leq t \leq t_{up}$ the stopping region has an upper r_t^u and a lower r_t^l threshold and, as long as S_t is between these two limits, we need to sample. Since the prior distribution is symmetric with respect to $1/2$, then, according to Theorem 1, the sampling region is symmetric around $t/2$, implying that $r_t^u + r_t^l = t$.

In Fig. 2, after using (24), we plot the average sample size

TABLE I
CHOICES OF k AND γ IN T_F IN (31) FOR 90%, 95%, AND 99% CONFIDENCE INTERVALS OF FIXED HALF-WIDTH h IN [20].

	90%		95%		99%	
h	k	γ	k	γ	k	γ
0.10	4	0.0754	4	0.0356	6	0.0068
0.05	4	0.0859	6	0.0433	8	0.0083
0.01	8	0.0972	10	0.0487	14	0.0097

and the two limits t_{lo}, t_{up} of T_o as functions of c for $a = 1$ and $h = 0.05$. We can see that the lower limit t_{lo} is significantly smaller than the resulting average, suggesting that the optimum scheme very quickly regards the accumulated information as capable of providing reliable interval estimates and therefore starts the process of questioning whether to stop or continue sampling.

IV. COMPARISONS

Let us now compare our scheme with the optimal fixed-sample-size (FSS) and two sequential methods: The first was proposed by Frey in [20] and the second, the *Conditional Method*, was proposed in our earlier work in [26]. Frey's method uses a modified Wald-type sequential confidence interval based on the stopping time

$$T_F = \inf \left\{ t \geq 0 : \frac{\tilde{\theta}_{t,k}(1 - \tilde{\theta}_{t,k})}{t} \leq \left(\frac{h}{z_{\gamma/2}} \right)^2 \right\}, \quad (31)$$

where $\tilde{\theta}_{t,k} = (S_t + k)/(t + 2k)$, $k > 0$ is a pre-specified constant and $\gamma = \gamma(k, h, \alpha)$ is chosen so that the confidence interval $[\hat{\theta}_{T_F} - h, \hat{\theta}_{T_F} + h]$, with $\hat{\theta}_t = S_t/t$, has a confidence level of at least $1 - \alpha$. Table I provides the values of k and γ recommended in [20] for best results.

From (31) and using the fact that $x(1 - x) \leq 1/4$ we conclude that the corresponding stopping time satisfies $T_F \leq \lceil \frac{1}{4h^2} z_{\gamma/2}^2 \rceil = N$. Regarding the finite-sample-size method, it uses the optimum Bayes estimator $\hat{\theta}_t$, obtained in (7) and the number of samples t is selected to meet the desired coverage probability. Finally, for the conditional method in [26], we should point out that it is a general sequential parameter estimation technique based on conditional costs which is not limited to binomial proportions. For the problem of interest, we have $T_C = \inf\{t \geq 0 : \mathcal{C}_t \leq \beta\}$ and $\hat{\theta}_{T_C} = \hat{\theta}_{T_C}$, where $\hat{\theta}_t, \mathcal{C}_t$ are the Bayes estimator and the corresponding optimum conditional complementary coverage probability defined in (7), (8). Threshold β is selected to guarantee that the resulting coverage probability is $1 - \alpha$. For \mathcal{C}_t we have from the proof of Theorem 1, eq. (33), that $\mathcal{C}_t \leq 2 \exp\{-2h^2(t + 2a + 1)\}$, consequently $T_C \leq \lceil \max\{\frac{1}{2h^2} |\log(\frac{\beta}{2})| - 2a - 1, 0\} \rceil = N$. In other words, all four schemes satisfy the assumption of Lemma 4 of bounded stopping time, therefore the corresponding performance can be computed numerically by applying the recursions of the lemma without the need to perform Monte-Carlo simulations.

For the competing methods using (24), (26), we plot in Fig. 3 the average number of samples $E[T]$ versus the coverage probability $P(|\hat{\theta}_T - \theta| \leq h)$ when $a = 1$ and $h = 0.05$.

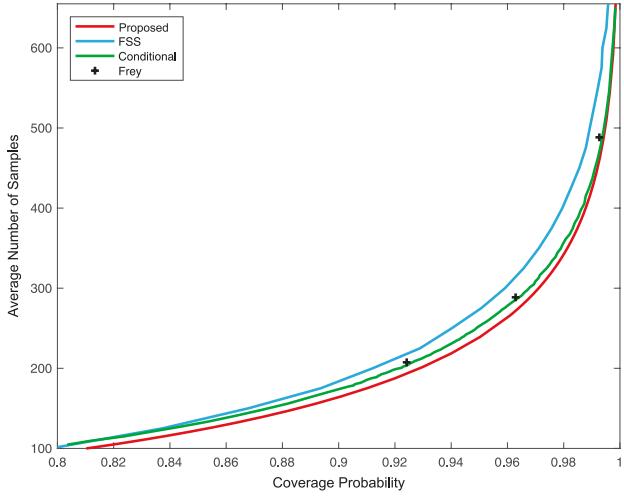


Fig. 3. Average samples size versus coverage probability for proposed (red), Frey (black +), fixed-sample-size (blue) and conditional (green), for $a = 1$ and $h = 0.05$.

Note that we have three points for Frey's scheme because of the tuning parameters k and γ which are provided in Table I only for three confidence levels. As we can see, the proposed method outperforms the fixed-sample-size and both alternative sequential techniques. It is only at very high coverage probability levels that the difference between the three sequential schemes becomes less pronounced.

As we pointed out in Remark 1, Section II, there is practical interest in evaluating the performance for each individual θ . Clearly in this case, the requirement is to be able to guarantee a minimal coverage probability for all θ . Again, we resort to Lemma 4 and use (23), (25) to evaluate the performance of the competing methods for each θ . In Fig. 4a, we plot the coverage probability for each test versus θ and in Fig. 4b, the corresponding average sample size required to obtain this performance. Parameters were selected so as all competing schemes provide the same worst-case coverage probability assuring a coverage of at least 0.95 for all θ . By observing the two figures, we can draw the following conclusions: The fixed-sample-size scheme can require up to almost eight times more samples compared to the proposed. Of course, one may argue that it produces higher coverage probability levels. Indeed this is true, but, unfortunately, this increased performance cannot be traded for a reduced sample size without compromising the worst-case level. Consequently, what we observe is in fact the best the fixed-sample-size method can offer. The conditional scheme, around $\theta = 0.5$, requires up to 30% more samples which, as in the case of fixed-sample-size, produce higher coverage probabilities. Again, it is impossible to sacrifice part of this increased performance to improve the corresponding sample size without degrading the worst-case coverage probability. Finally, we can see that the proposed and Frey's scheme require similar samples over most θ . However, we observe that the proposed method has a coverage probability profile which is better than Frey's, since for most θ the corresponding probability is larger. Frey's scheme is slightly better only for θ close to 0 and 1. But even for these values of θ the proposed

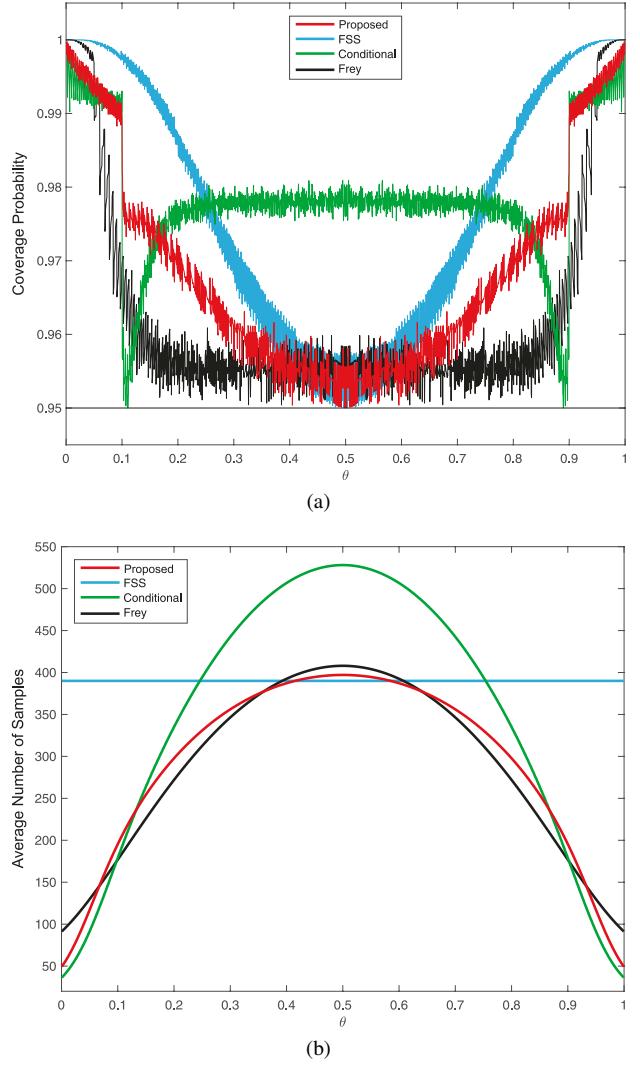


Fig. 4. Coverage probability (a) and Average sample size (b) as a function of proportion θ for proposed (red), Frey (black), fixed-sample-size (blue) and conditional (green) when $a = 1$, $h = 0.05$ and worst-case coverage probability 0.95.

scheme requires almost 50% less samples.

V. CONCLUSIONS

We proposed an optimal sequential scheme for obtaining confidence intervals for a binomial proportion under a well defined formulation. We proved that, for a particular prior (Beta density), our optimum stopping time enjoys certain uncommon properties not encountered in solutions of other classical optimal stopping problems. We also conjectured that these properties are present with any prior. Specifically, our claim is that our stopping time is always bounded from above and below, suggesting that we need to first accumulate a sufficient amount of information before we start applying our stopping rule, and that our stopping time will always terminate before or at a specific deterministic time even if we allow the time horizon to be infinite. Finally, our scheme was compared against the optimum fixed-sample-size procedure and against existing sequential alternatives. Numerical performance eval-

ulations showed that the proposed method exhibits an overall improved performance profile compared to its rivals.

VI. ACKNOWLEDGMENTS

The authors are grateful to the Associate Editor and the two anonymous reviewers for their constructive comments that greatly improved the quality and presentation of this article.

APPENDIX

Proof of Lemma 1: When $E[T] = \infty$, then (10) trivially holds, since both sides are equal to infinity. Thus, it is sufficient to consider the case when $E[T] < \infty$, which implies that T is finite with probability (w.p.) 1. For any given integer N , from (10) we have

$$\begin{aligned} J(T, \hat{\theta}_T) &= cE[T] + P(|\hat{\theta}_T - \theta| > h) \\ &= E \left[\sum_{t=0}^{\infty} \left\{ ct + \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \right\} \mathbb{1}_{\{T=t\}} \right] \\ &= \lim_{N \rightarrow \infty} E \left[\sum_{t=0}^N \left\{ ct + \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \right\} \mathbb{1}_{\{T=t\}} \right] \end{aligned} \quad (32)$$

$$\begin{aligned} &= \lim_{N \rightarrow \infty} \sum_{t=0}^N E \left[\left\{ ct + \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \right\} \mathbb{1}_{\{T=t\}} \right] \\ &= \lim_{N \rightarrow \infty} \sum_{t=0}^N E \left[\left\{ ct + P(|\hat{\theta}_t - \theta| > h | \mathcal{F}_t) \right\} \mathbb{1}_{\{T=t\}} \right] \end{aligned} \quad (33)$$

$$\begin{aligned} &\geq \lim_{N \rightarrow \infty} \sum_{t=0}^N E \left[\left\{ ct + \inf_{\hat{\theta}_t} P(|\hat{\theta}_t - \theta| > h | \mathcal{F}_t) \right\} \mathbb{1}_{\{T=t\}} \right] \end{aligned} \quad (34)$$

$$\begin{aligned} &= \lim_{N \rightarrow \infty} \sum_{t=0}^N E \left[\{ct + \mathcal{C}_t\} \mathbb{1}_{\{T=t\}} \right] \\ &= \lim_{N \rightarrow \infty} E \left[\sum_{t=0}^N \{ct + \mathcal{C}_t\} \mathbb{1}_{\{T=t\}} \right] \\ &= E \left[\sum_{t=0}^{\infty} \{ct + \mathcal{C}_t\} \mathbb{1}_{\{T=t\}} \right] = E[ct + \mathcal{C}_T]. \end{aligned} \quad (35)$$

The exchange of limit and expectation in (32) is possible, since the sum contains nonnegative terms and is therefore monotone increasing in N with a limit having finite expectation. Hence monotone convergence assures the validity of the exchange. Equation (33) is true because $\mathbb{1}_{\{T=t\}}$ is \mathcal{F}_t -measurable, while we have equality in (34) if we select $\hat{\theta}_t = \hat{\theta}_t$ when $\{T=t\}$. Finally the exchange of limit and expectation in (35) is again possible due to monotone convergence. This completes the proof. \square

Proof of Theorem 1: If $\alpha \geq P(|\hat{\theta}_0 - \theta| > h)$ then stopping at $T_o = 0$ corresponds to the smallest possible (average) number of samples while, at the same time, we satisfy the coverage probability constraint.

To prove ii) we first show that there exists N_α such that $P(|\hat{\theta}_{N_\alpha} - \theta| > h) < \alpha$. Note that

$$\begin{aligned} P(|\hat{\theta}_t - \theta| > h) &\leq P \left(\left| \frac{S_t}{t} - \theta \right| > h \right) \\ &\leq \frac{1}{h^2} E \left[\left(\frac{S_t}{t} - \theta \right)^2 \right] \\ &= \frac{1}{h^2} E \left[E \left[\left(\frac{S_t}{t} - \theta \right)^2 | \theta \right] \right] \\ &= \frac{1}{h^2} E \left[\frac{\theta(1-\theta)}{t} \right] \\ &\leq \frac{1}{4h^2 t}, \end{aligned} \quad (36)$$

where we used the fact that S_t/t is not the optimum Bayes estimator of the mid-point, then we applied the Chebyshev's inequality, then the fact that S_t/t is an estimator of θ with estimation error variance equal to $\theta(1-\theta)/t$ and finally that $\theta(1-\theta) \leq 1/4$. From (36) we conclude that $P(|\hat{\theta}_t - \theta| > h) \rightarrow 0$ as $t \rightarrow \infty$ therefore, there exists N_α such that $P(|\hat{\theta}_{N_\alpha} - \theta| > h) < \alpha$.

Fix $N \geq N_\alpha$ and denote $\mathcal{V}_t(S_t, c) = \inf_{t \leq T \leq N} E[c(T-t) + \mathcal{C}_T | \mathcal{F}_t]$, where we underline the dependence of \mathcal{V}_t on c (in addition to S_t). For $0 \leq c_1 \leq c_2$ and $T \geq t$ we can write

$$c_1(T-t) + \mathcal{C}_T \leq c_2(T-t) + \mathcal{C}_T,$$

which, after taking expectation conditioned on \mathcal{F}_t and then infimum over $t \leq T \leq N$, proves that $\mathcal{V}_t(S_t, c)$ is increasing in c . The increase of $\mathcal{V}_t(S_t, c)$ with respect to c also suggests that the optimum stopping time $T_o(c)$, defined in (17), is a decreasing function of c .

Consider now the sequence of optimum complementary coverage probabilities $\{\mathcal{C}_t\}$, we observe

$$\begin{aligned} \mathcal{C}_t &= \inf_{\hat{\theta}} P(|\hat{\theta} - \theta| > h | \mathcal{F}_t) = \inf_{\hat{\theta}} E[P(|\hat{\theta} - \theta| > h | \mathcal{F}_{t+1}) | \mathcal{F}_t] \\ &\geq E \left[\inf_{\hat{\theta}} P(|\hat{\theta} - \theta| > h | \mathcal{F}_{t+1}) | \mathcal{F}_t \right] = E[\mathcal{C}_{t+1} | \mathcal{F}_t]. \end{aligned} \quad (37)$$

We can then write

$$\begin{aligned} P(|\hat{\theta}_{T_o(c)} - \theta| > h) &= E[\mathcal{C}_{T_o(c)}] \\ &= \mathcal{C}_0 - E \left[\sum_{t=0}^{T_o(c)-1} \{\mathcal{C}_t - \mathcal{C}_{t+1}\} \right] \\ &= \mathcal{C}_0 - E \left[\sum_{t=0}^N \{\mathcal{C}_t - \mathcal{C}_{t+1}\} \mathbb{1}_{\{T_o(c) > t\}} \right] \\ &= \mathcal{C}_0 - E \left[\sum_{t=0}^N \{\mathcal{C}_t - E[\mathcal{C}_{t+1} | \mathcal{F}_t]\} \mathbb{1}_{\{T_o(c) > t\}} \right], \end{aligned} \quad (38)$$

where for the last equality we used the fact that $\mathbb{1}_{\{T_o(c) > t\}}$ is \mathcal{F}_t -measurable. This combined with (37) and the decrease of $T_o(c)$ with respect to c , implies that $P(|\hat{\theta}_{T_o(c)} - \theta| > h)$ is increasing in c .

For $c = 1$ we stop at 0 and, therefore, $P(|\hat{\theta}_{T_o(1)} - \theta| > h) = P(|\hat{\theta}_0 - \theta| > h) > \alpha$. Set now $c = 0$ which suggests that the cost of sampling is zero and therefore the optimum is to stop

at N (we also deduce this by combining (17) and (37)). This yields $\mathbb{P}(|\hat{\theta}_{T_o(0)} - \theta| > h) = \mathbb{P}(|\hat{\theta}_N - \theta| > h) = \mathbb{E}[\mathcal{C}_N]$. Now from (37) by averaging we conclude that $\mathbb{E}[\mathcal{C}_t]$ is decreasing in t and for $N > N_\alpha$ we have $\mathbb{E}[\mathcal{C}_N] \leq \mathbb{E}[\mathcal{C}_{N_\alpha}] < \alpha$, implying $\mathbb{P}(|\hat{\theta}_{T_o(0)} - \theta| > h) < \alpha$. As mentioned, $\mathbb{P}(|\hat{\theta}_{T_o(c)} - \theta| > h)$ is increasing in c , if it is also continuous then there exists $0 < c_* < 1$ satisfying $\mathbb{P}(|\hat{\theta}_{T_o(c_*)} - \theta| > h) = \alpha$ which means that $T_o(c_*)$ solves the constrained problem. In case the function $\mathbb{P}(|\hat{\theta}_{T_o(c)} - \theta| > h)$ exhibits a jump at c_* such that for c_*- the probability is strictly smaller than α while for c_*+ it is strictly larger, then before taking any samples we need to perform a randomization to decide which of the two stopping times $T_o(c_*-), T_o(c_*)$ to use. The randomization probability must be selected so that we satisfy the constraint with equality. \square

Proof of Lemma 4: We prove (23) first. Set $\Omega_N = \emptyset$, i.e. we stop necessarily at N . Then we note that

$$T = \sum_{t=0}^{N-1} \mathbb{1}_{\{T > t\}} = (1 + \cdots (1 + \mathbb{1}_{\{S_{N-1} \in \Omega_{N-2}\}} (1 + \mathbb{1}_{\{S_{N-1} \in \Omega_{N-1}\}} (1 + \mathbb{1}_{\{S_N \in \Omega_N\}}))) \cdots)$$

suggesting that

$$\mathbb{E}[T|\theta] = \mathbb{E}[(1 + \cdots \mathbb{E}[(1 + \mathbb{1}_{\{S_{N-1} \in \Omega_{N-1}\}} \mathbb{E}[(1 + \mathbb{1}_{\{S_N \in \Omega_N\}}) | \mathcal{F}_{N-1}, \theta] | \mathcal{F}_{N-2}, \theta] \cdots) | \theta].$$

If we set $U_N(S_N) = 0$ then we can define the backward recursion

$$\begin{aligned} U_t(S_t) &= \mathbb{E}[1 + \mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_{t+1}) | \mathcal{F}_t] \\ &= 1 + \mathbb{E}[\mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_{t+1}) | \mathcal{F}_t] \\ &= 1 + \mathbb{P}(X_{t+1} = 1 | S_t, \theta) \mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_t + 1) \\ &\quad + \mathbb{P}(X_{t+1} = 0 | S_t, \theta) \mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_t) \\ &= 1 + \theta \mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_t + 1) + (1 - \theta) \mathbb{1}_{\{S_{t+1} \in \Omega_{t+1}\}} U_{t+1}(S_t), \end{aligned}$$

which proves (23) and, also, that $U_0(S_0) = \mathbb{E}[T|\theta]$. For (24) we proceed similarly the only difference being that $\mathbb{P}(X_{t+1} = 1 | S_t) = g_{t+1}(S_t)$ with this probability being defined in (16).

For (25) and (26) we follow similar steps. We have

$$\begin{aligned} \mathbb{1}_{\{|\hat{\theta}_T - \theta| > h\}} &= \sum_{t=0}^N \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \mathbb{1}_{\{T=t\}} \\ &= \sum_{t=0}^N \mathbb{1}_{\{|\hat{\theta}_t - \theta| > h\}} \mathbb{1}_{\{S_t \notin \Omega_t\}} \prod_{j=0}^{t-1} \mathbb{1}_{\{S_j \in \Omega_j\}} \\ &= (\mathbb{1}_{\{|\hat{\theta}_0 - \theta| > h\}} \mathbb{1}_{\{S_0 \notin \Omega_0\}}) + (\mathbb{1}_{\{|\hat{\theta}_1 - \theta| > h\}} \mathbb{1}_{\{S_1 \notin \Omega_1\}}) \mathbb{1}_{\{S_0 \in \Omega_0\}} \\ &\quad + \cdots + (\mathbb{1}_{\{|\hat{\theta}_N - \theta| > h\}} \mathbb{1}_{\{S_N \notin \Omega_N\}}) \prod_{j=1}^{N-1} \mathbb{1}_{\{S_j \in \Omega_j\}}. \end{aligned}$$

Applying expectation given θ yields

$$\begin{aligned} \mathbb{P}(|\hat{\theta}_T - \theta| > h | \theta) &= \mathbb{E}[\mathbb{1}_{\{|\hat{\theta}_0 - \theta| > h\}} \mathbb{1}_{\{S_0 \notin \Omega_0\}} + \cdots \\ &\quad + \mathbb{E}[\mathbb{1}_{\{|\hat{\theta}_{N-1} - \theta| > h\}} \mathbb{1}_{\{S_{N-1} \notin \Omega_{N-1}\}}] \\ &\quad + (\mathbb{E}[\mathbb{1}_{\{|\hat{\theta}_N - \theta| > h\}} \mathbb{1}_{\{S_N \notin \Omega_N\}} | \mathcal{F}_{N-1}, \theta]) \cdots | \theta]. \end{aligned}$$

Defining $W_N(S_N) = \mathbb{1}_{\{|\hat{\theta}_N - \theta| > h\}}$ it is straightforward to see that the recursion in (25) computes the desired complementary coverage probability. Similarly for (26) only now instead of conditioning with respect to both \mathcal{F}_t and θ we condition only with respect to \mathcal{F}_t . This concludes the proof. \square

Proof of Theorem 2: Let us first find upper and lower bounds of $\mathcal{C}_t(S_t)$ that are independent from S_t . From [27, Theorem 2.1] and for a random variable X with density $\text{Beta}(x, p, q)$ we have that

$$\mathbb{E}[e^{\lambda(X - \mu)}] \leq e^{\frac{\lambda^2}{8(p+q+1)}}, \quad \lambda > 0, \quad (39)$$

where $\mu = \frac{p}{p+q}$ is the average under the Beta density. Using the Markov inequality we can then write

$$\begin{aligned} \mathbb{P}(|X - \mu| > h) &= \mathbb{P}(X - \mu > h) + \mathbb{P}(X - \mu < -h) \\ &= \mathbb{P}(X - \mu > h) + \mathbb{P}(1 - X - (1 - \mu) > h) \\ &\leq \frac{\mathbb{E}[e^{\lambda(X - \mu)}]}{e^{\lambda h}} + \frac{\mathbb{E}[e^{\lambda(1 - X - (1 - \mu))}]}{e^{\lambda h}} \leq 2e^{\frac{\lambda^2}{8(p+q+1)} - \lambda h}, \end{aligned} \quad (40)$$

where we used the fact that if X is Beta distributed with parameters p, q then $1 - X$ is also Beta with parameters q, p . Selecting in (40) $\lambda = 4(p+q+1)h$ yields the tightest upper bound, namely

$$\mathbb{P}(|X - \mu| > h) \leq 2e^{-2h^2(p+q+1)}. \quad (41)$$

We can now use this result to upper bound $\mathcal{C}_t(S_t)$. We observe that

$$\begin{aligned} \mathcal{C}_t(S_t) &= \inf_{\hat{\theta}_t} \mathbb{P}(|\theta - \hat{\theta}_t| > h | \mathcal{F}_t) \\ &\leq \mathbb{P}(|\theta - \mathbb{E}[\theta | \mathcal{F}_t]| > h | \mathcal{F}_t) \leq 2e^{-2h^2(t+2a+1)}. \end{aligned} \quad (42)$$

For the last inequality we used (41) and the fact that θ given \mathcal{F}_t is Beta distributed with parameters $p = S_t + a$ and $q = t - S_t + a$.

Let us now find a lower bound for $\mathcal{C}_t(S_t)$. From [25, Page 944, Formula 26.5.15] we conclude that $I_x(p, q) > I_x(p+1, q-1)$ for $q > 1$. Using this inequality repeatedly in (28) we conclude

$$\begin{aligned} \mathbb{P}(|\hat{\theta}_t - \theta| > h | \mathcal{F}_t) &= 1 - I_{\min\{1, \hat{\theta}_t + h\}}(S_t + a, t - S_t + a) \\ &\quad + I_{\max(0, \hat{\theta}_t - h)}(S_t + a, t - S_t + a) \\ &= I_{\max\{0, 1 - h - \hat{\theta}_t\}}(t - S_t + a, S_t + a) \\ &\quad + I_{\max(0, \hat{\theta}_t - h)}(S_t + a, t - S_t + a) \\ &\geq I_{\max\{0, 1 - h - \hat{\theta}_t\}}(t + 2n_a + \delta_a, \delta_a) \\ &\quad + I_{\max\{0, \hat{\theta}_t - h\}}(t + 2n_a + \delta_a, \delta_a), \end{aligned} \quad (43)$$

where for the second equality we used the property $1 - I_x(p, q) = I_{1-x}(q, p)$ and where n_a, δ_a are defined as

$$n_a = \begin{cases} [a] & \text{if } a \text{ not an integer} \\ a - 1 & \text{if } a \text{ an integer,} \end{cases}$$

and

$$\delta_a = \begin{cases} a - [a] & \text{if } a \text{ not an integer} \\ 1 & \text{if } a \text{ an integer,} \end{cases}$$

where $[a]$ denotes integer part of a . Since $a > 0$ we have $n_a \geq 0, 1 \geq \delta_a > 0$ and $a = n_a + \delta_a$. By taking the derivative

of the last sum in (43) with respect to $\hat{\theta}_t$ we can show that it has the same sign as the following expression

$$\phi(\hat{\theta}_t) = \frac{(\hat{\theta}_t - h)^{t+2n_a + \delta_a - 1}}{(1 + h - \hat{\theta}_t)^{1-\delta_a}} - \frac{(1 - h - \hat{\theta}_t)^{t+2n_a + \delta_a - 1}}{(\hat{\theta}_t + h)^{1-\delta_a}}.$$

Now it is easy to verify that $\phi(1 - \hat{\theta}_t) = -\phi(\hat{\theta}_t)$ therefore it is sufficient to analyze the sign of $\phi(\hat{\theta}_t)$ for $h \leq \hat{\theta}_t \leq 0.5$. When $t \geq 1$ and because $1 \geq \delta_a$ we can see that the sign is negative for any value of a , suggesting that we have a minimum for $\hat{\theta}_t = 0.5$. Therefore, if $\Gamma(x)$ denotes the Gamma function, then for $t \geq 1$ we can write

$$\begin{aligned} \mathcal{C}_t &\geq 2I_{0.5-h}(t+2n_a + \delta_a, \delta_a) \\ &\geq 2 \frac{\Gamma(t+2n_a + 2\delta_a)(0.25 - h^2)^{\delta_a}}{\Gamma(t+2n_a + \delta_a + 1)\Gamma(\delta_a)} (0.5 - h)^{t+2n_a} \\ &= 2 \frac{\Gamma(t+2n_a + 2\delta_a + 1)(0.25 - h^2)^{\delta_a}}{(t+2n_a + 2\delta_a)\Gamma(t+2n_a + \delta_a + 1)\Gamma(\delta_a)} (0.5 - h)^{t+2n_a} \\ &\geq 2 \frac{(0.25 - h^2)^{\delta_a}}{(t+2n_a + 2\delta_a)\Gamma(\delta_a)} (0.5 - h)^{t+2n_a}. \end{aligned} \quad (44)$$

In the previous expression the second inequality comes from [25, Page 944, Formula 26.5.16]; for the next equality we used the property $\Gamma(x+1) = x\Gamma(x)$; while for the last inequality we used the increase of $\Gamma(x)$ for $x \geq 1.5$, which is true in our case for $t \geq 1$ and any $a > 0$.

Having established bounds for \mathcal{C}_t we can now compute an upper bound N for t_{up} and a lower bound ν for t_{lo} therefore proving their existence and demonstrating properties i) and ii). We first note that if $\mathcal{C}_N \leq c$ in (14) we will have $\mathcal{C}_N \leq c + \mathcal{V}_N$ meaning that $\mathcal{V}_N = \mathcal{C}_N$ and consequently N is a stopping instant for *all* values of S_t . This implies that $T_{\text{o}} \leq N$. Quantity t_{up} is the smallest N for which this inequality is true for all S_t . Requiring $2e^{-2h^2(N+2a+1)} \leq c$ we obtain

$$N = \left\lceil \max \left\{ 0, \frac{|\log(c)| + \log(2)}{2h^2} - 2a - 1 \right\} \right\rceil. \quad (45)$$

To find a lower bound ν for t_{lo} we combine the lower bound of \mathcal{C}_t with an upper bound for \mathcal{V}_t . Finding the latter is straightforward. Indeed if we start from time instant N which, as we argued, is selected so that $\mathcal{C}_N \leq c$, then using induction and the fact that

$$\mathcal{V}_t = \min\{\mathcal{C}_t, c + \mathbb{E}[V_{t+1}|\mathcal{F}_t]\} \leq c + \mathbb{E}[\mathcal{V}_{t+1}|\mathcal{F}_t]$$

we can show that $\mathcal{V}_t \leq c + c(N-t) = c(N+1-t)$. It is then clear that, as long as $c(N+1) \leq \mathcal{C}_0$, for any $t \geq 1$ for which we have

$$c(N+1-t)(t+2n_a + 2\delta_a) \leq 2 \frac{(0.25 - h^2)^{\delta_a}}{\Gamma(\delta_a)} (0.5 - h)^{t+2n_a} \quad (46)$$

we do not stop at this time instant. In fact we can see that we have an interval of the form $t \in [0, \dots, \nu]$ during which no stopping can occur. A rough estimate of ν can be obtained by solving instead of (46) the simpler alternative

$$\begin{aligned} \max_t c(N+1)(t+2n_a + 2\delta_a) &= \frac{c}{4}(N+2n_a + 2\delta_a)^2 \\ &\leq \frac{2(0.25 - h^2)^{\delta_a}}{\Gamma(\delta_a)} (0.5 - h)^{\nu+2n_a}, \end{aligned}$$

which yields

$$\nu = \left\lceil \max \left\{ 0, \frac{|\log(c)| - \log(I_1) + \log(I_2)}{|\log(0.5 - h)|} - 2n_a \right\} \right\rceil, \quad (47)$$

where

$$I_1 = (N+2n_a + \delta_a)^2 \Gamma(\delta_a), \quad I_2 = 8(0.25 - h^2)^{\delta_a},$$

provided c satisfies $c \leq \frac{\mathcal{C}_0}{N+1}$. Regarding the latter, if we are in the non-trivial case where we do not stop at time 0 then $\alpha < \mathcal{C}_0$, consequently it is sufficient to have $c \leq \frac{\alpha}{N+1}$. We thus conclude that for small enough c there is a lower limit $t_{\text{lo}} \geq \nu$ which is nontrivial. This concludes the proof. \square

REFERENCES

- [1] A. Sullivan, D. Raben, J. Reekie, M. Rayment, A. Mocroft *et al.*, “Feasibility and effectiveness of indicator condition-guided testing for HIV: results from HIDES I (HIV indicator diseases across Europe study),” *PLoS ONE*, vol. 8, no. 1, p. e52845, 2013.
- [2] J. Abramson, R. Takvorian, D. Fisher, Y. Feng, E. Jacobsen *et al.*, “Oral clofarabine for relapsed/refractory non-Hodgkin lymphomas: results of a phase 1 study,” *Leukemia & Lymphoma*, vol. 54, pp. 1915–1920, 2013.
- [3] J. T. Morisette and S. Khorram, “Exact binomial confidence interval for proportions,” *Photogrammetric Engineering & Remote Sensing*, vol. 64, no. 4, pp. 281–283, 1998.
- [4] A. Agresti and B. A. Coull, “Approximate is better than ‘exact’ for interval estimation of binomial proportions,” *The American Statistician*, vol. 52, pp. 119–126, 1998.
- [5] L. D. Brown, T. Cai, and A. Dasgupta, “Interval estimation for a binomial proportion (with discussion),” *Statistical Science*, vol. 16, pp. 101–133, 2001.
- [6] R. G. Newcombe, “Two-sided confidence intervals for the single proportion: comparison of seven methods,” *Statistics in Medicine*, vol. 17, pp. 857–872, 1998.
- [7] S. E. Vollset, “Confidence intervals for a binomial proportion,” *Statistics in Medicine*, vol. 12, no. 9, pp. 809–824, 1993.
- [8] E. B. Wilson, “Probable inference, the law of succession, and statistical inference,” *Journal of the American Statistical Association*, vol. 22, pp. 209–212, 1927.
- [9] C. J. Clopper and E. S. Pearson, “The use of confidence or fiducial limits illustrated in the case of the binomial,” *Biometrika*, vol. 26, pp. 404–413, 1934.
- [10] T. E. Sterne, “Some remarks on confidence of fiducial limits,” *Biometrika*, vol. 41, no. 1, pp. 275–278, 1954.
- [11] E. L. Crow, “Confidence intervals for a proportion,” *Biometrika*, vol. 43, pp. 423–435, 1956.
- [12] C. R. Blyth and H. A. Still, “Binomial confidence intervals,” *Journal of the American Statistical Association*, vol. 78, pp. 108–116, 1983.
- [13] L. D. Brown, T. Cai, and A. DasGupta, “Interval estimation for a binomial proportion and asymptotic expansions,” *The Annals of Statistics*, vol. 30, pp. 160–201, 2002.
- [14] J. Reiczigel, “Confidence intervals for the binomial parameter: some new considerations,” *Statistics in Medicine*, vol. 22, no. 4, pp. 611–621, 2003.
- [15] A. M. Pires and C. Amado, “Interval estimators for a binomial proportion: comparison of twenty methods,” *REVSTAT - Statistical Journal*, vol. 6, no. 2, pp. 165–197, 2008.
- [16] M. Thulin, “The cost of using exact confidence intervals for a binomial proportion,” *Electron. J. Statist.*, vol. 8, no. 1, pp. 817–840, 2014.
- [17] C. Jegourel, J. Sun, and J. S. Dong, *Sequential Schemes for Frequentist Estimation of Properties in Statistical Model Checking*. Springer International Publishing, 2017, pp. 333–350.
- [18] Y. S. Chow and H. Robbins, “On the asymptotic theory of fixed-width sequential confidence intervals for the mean,” *The Annals of Mathematical Statistics*, vol. 36, no. 2, pp. 457–462, 1965.
- [19] M. Tanaka, “On a confidence interval of given length for the parameter of the binomial and the Poisson distributions,” *Annals of the Institute of Statistical Mathematics*, vol. 13, pp. 201–215, 1961.
- [20] J. Frey, “Fixed-width sequential confidence intervals for a proportion,” *The American Statistician*, vol. 64, pp. 242–249, 2010.
- [21] P. Cabilio, “Sequential estimation in Bernoulli trials,” *The Annals of Statistics*, vol. 5, no. 2, pp. 342–356, 1977.

- [22] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer, 1988.
- [23] A. N. Shiryaev, *Optimal Stopping Rules*. Springer, 1978.
- [24] T. Choi and R. V. Ramamoorthi, *Remarks on consistency of posterior distributions*, ser. Collections. Beachwood, Ohio, USA: Institute of Mathematical Statistics, 2008, vol. Volume 3, pp. 170–186.
- [25] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*, 9th ed. Dover Publication, 1972.
- [26] G. V. Moustakides, T. Yaacoub, and Y. Mei, “Sequential estimation based on conditional cost,” *Proceedings of IEEE International Symposium on Information Theory*, pp. 436–440, June 2017.
- [27] O. Marchal and J. Arbel, “On the sub-Gaussianity of the Beta and Dirichlet distributions,” *Electron. Commun. Probab.*, vol. 22, no. 54, pp. 1–14, 2017.

Tony Yaacoub received his Ph.D. degree in industrial engineering from the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology, Atlanta, GA, in 2018. He is currently with the pricing optimization team at Delta Air Lines, Inc. His research interests include sequential detection and estimation, verification and validation of simulation models, and machine learning.

George V. Moustakides received the diploma in Electrical and Mechanical Engineering from the National Technical University of Athens, Greece in 1979, the MSE in Systems Engineering from the University of Pennsylvania in 1980, and the M.Sc and Ph.D in Electrical Engineering and Computer Science from Princeton University in 1983. Since 1988 he is with the University of Patras, Greece initially with the Computer Engineering and Informatics department and afterwards with the department of Electrical and Computer Engineering. In 2017 he also joined the Computer Science department at Rutgers University. In the past he held various visiting or long-term appointments with INRIA, Princeton University, University of Pennsylvania, Columbia University, University of Maryland, Georgia Institute of Technology, University of Southern California, University of Illinois at Urbana-Champaign and Rutgers University. His interests include Sequential Detection, Statistical Signal Processing and Machine Learning. From 2011 to 2014 and from 2016 to 2018 he served as Associate Editor for the IEEE Transactions on Information Theory.

Yajun Mei received the B.S. degree in Mathematics from Peking University, Beijing, P. R. China in 1996, and the Ph. D. degree in Mathematics with a minor in Electrical Engineering in 2003 from the California Institute of Technology, Pasadena, CA, USA.

He is currently an Associate Professor in The H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology, Atlanta, GA. His research interests include streaming data analysis, quickest change detection, sensor networks, sequential analysis, and their applications in engineering and biomedical sciences.