Published in final edited form as:

J Stat Plan Inference. 2014 October 1; 153: 100–114. doi:10.1016/j.jspi.2014.05.010.

Sequential Tests of Multiple Hypotheses Controlling Type I and II Familywise Error Rates

Jay Bartroff* and Jinlin Song†

Jay Bartroff: bartroff@usc.edu; Jinlin Song: jsong@analysisgroup.com

*Department of Mathematics, University of Southern California, Los Angeles, California, USA

†Analysis Group, 111 Huntington Avenue, Tenth Floor, Boston, MA 02199, USA

Abstract

This paper addresses the following general scenario: A scientist wishes to perform a battery of experiments, each generating a sequential stream of data, to investigate some phenomenon. The scientist would like to control the overall error rate in order to draw statistically-valid conclusions from each experiment, while being as efficient as possible. The between-stream data may differ in distribution and dimension but also may be highly correlated, even duplicated exactly in some cases. Treating each experiment as a hypothesis test and adopting the familywise error rate (FWER) metric, we give a procedure that sequentially tests each hypothesis while controlling both the type I and II FWERs regardless of the between-stream correlation, and only requires arbitrary sequential test statistics that control the error rates for a given stream in isolation. The proposed procedure, which we call the sequential Holm procedure because of its inspiration from Holm's (1979) seminal fixed-sample procedure, shows simultaneous savings in expected sample size and less conservative error control relative to fixed sample, sequential Bonferroni, and other recently proposed sequential procedures in a simulation study.

Key words and phrases

Clinical trials; Holm's procedure; Multiple testing; Sequential analysis; Step-down test; Wald approximations

1 Introduction

This paper addresses the following scenario: A scientist wishes to perform a battery of k experiments sequentially in time in order to investigate some phenomenon, resulting in k data streams:

© 2014 Elsevier B.V. All rights reserved.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

$$\begin{array}{lll}
\text{Data stream 1} & X_1^{(1)}, X_2^{(1)}, \dots & \text{from Experiment 1} \\
\text{Data stream 2} & X_1^{(2)}, X_2^{(2)}, \dots & \text{from Experiment 2} \\
\vdots & & \\
\text{Data stream } k & X_1^{(k)}, X_2^{(k)}, \dots & \text{from Experiment } k.
\end{array} \tag{1}$$

The scientist would like to control the overall error rate of the battery of experiments in order to be able to draw statistically-valid conclusions for each experiment once all experimentation has ceased, but also needs to be as efficient as possible with the finite resources available by “dropping” certain experiments (i.e., stopping experimentation) when additional data is no longer needed from that stream to reach a conclusion. The between-stream data may be very dissimilar in distribution and dimension, but at the same time may be highly correlated, or even duplicated exactly in some cases, since they all are related to some phenomenon.

The preceding scenario occurs in a number of real applications including multiple endpoint (or multi-arm) clinical trials (Jennison and Turnbull, 2000, Chapter 15), multi-channel changepoint detection (Tartakovsky et al., 2003) and its applications to biosurveillance (Mei, 2010), genetics and genomics (Dudoit and van der Laan, 2008), acceptance sampling with multiple criteria (Baillie, 1987), and financial trading strategies (Romano and Wolf, 2005). If we think of each experiment as a hypothesis test about the corresponding data stream, then what is needed is a combination of a multiple hypothesis test and a sequential hypothesis test. We point out that our use of the word “sequential” here and below refers to the manner of sampling (or equivalently, observation) and differs from the way the word is sometimes used in the literature on fixed-sample multiple testing procedures to describe the stepwise analysis of fixed-sample test statistics, e.g., p -values.

This scenario described above was addressed by Bartroff and Lai (2010) who gave a procedure that sequentially (or group sequentially) tests k hypotheses while controlling the type I familywise error rate (FWER, see Hochberg and Tamhane, 1987), i.e., the probability of rejecting any true hypotheses, at a prescribed level. Their procedure requires only the existence of basic sequential tests for each data stream and makes no assumptions about the dependence between the different data streams; in particular, the error control holds when the streams are highly positively correlated, as is often the case in the application areas mentioned above. The current paper introduces a procedure to test k hypotheses while simultaneously controlling both the type I and II FWERs (defined precisely below) at prescribed levels in the same general setting: No assumptions are made about the dependence between the different data streams. We call this new procedure the *sequential Holm procedure* because of its relation to Holm’s (1979) seminal fixed-sample “step-down” procedure which controls the FWER. Following a review of relevant previous work, we give a general formulation of the sequential Holm procedure in Section 2. In Section 3 we consider simple hypotheses and then composite hypotheses, simulation studies in Section 4, and finally a discussion of future extensions and a summary.

1.1 Background and Previous Work

Separately, multiple testing and sequential testing are both quite mature fields, the latter dating back to Wald's (1947) invention of sequential analysis following World War II (see Siegmund (1985) for a summary of the major developments). Work on multiple testing dates back to classical "multiple comparison" procedures of Fisher (1932), Scheffé (1953), Tukey, and others (see Seber and Lee, 2003) for testing hypotheses about parameter vectors in linear models. Holm (1979) proposed a general method for testing a list of null hypotheses $H^{(1)}, \dots, H^{(k)}$ that controls the type I FWER without making any assumptions about the structure of the hypotheses or correlations between the valid, fixed-sample p -values $p^{(1)}, \dots, p^{(k)}$ associated with $H^{(1)}, \dots, H^{(k)}$, respectively. By relabeling if necessary, without loss of generality let $p^{(1)} \leq \dots \leq p^{(k)}$. Given a prescribed FWER bound $\alpha \in (0, 1)$, at stage $j = 1, \dots, k$, Holm's procedure accepts $H^{(j)}, \dots, H^{(k)}$ and terminates if $p^{(j)} \leq \alpha/(k - j + 1)$. Otherwise, $H^{(j)}$ is rejected and stage $j + 1$ is commenced (provided $j < k$). Lehmann and Romano (2005) give a simple proof that the FWER of Holm's procedure is bounded by α .

The intersection of multiple testing and sequential testing is less well-developed in a general setting than either individual field. One area that has been considered is the adaptation of some classical fixed-sample tests about vector parameters, such as those mentioned above, to the sequential sampling setting, including O'Brien and Fleming's (1979) sequential version of Pearson's χ^2 test, and Tang et al.'s (1989, 1993) group sequential extensions of O'Brien's (1984) generalized least squares statistic. For bivariate normal populations, Jennison and Turnbull (1993) proposed a sequential test of two one-sided hypotheses about the bivariate mean vector, and Cook and Farewell (1994) proposed a sequential test in a similar setting but where one of the hypotheses is two-sided. A procedure for comparing three treatments was proposed by Siegmund (1993), related to Paulson's (1964) earlier procedure for selecting the largest mean of k normal distributions, which Bartroff and Lai (2010) showed to be a special case of their more general sequential step-down method, mentioned in the previous paragraph. Recently, Ye et al. (2013) rediscovered another special case of Bartroff and Lai's (2010) procedure.

All of these procedures mentioned so far aim to explicitly control either the classical type I error probability or the type I FWER. The first sequential procedures to simultaneously control both the type I and II FWERs were introduced by De and Baron (2012a,b), who consider the n th set of k measurements

$$X_n^{(1)}, X_n^{(2)}, \dots, X_n^{(k)} \quad (2)$$

from the streams (1) to be the measurements on the n th patient in a study, and patients are sampled until accept/reject decisions can be reached for each data stream. These procedures are compared with the proposed sequential Holm procedure in the numerical comparisons in Section 4. The need to drop certain measurements (or endpoints), even on a patient remaining in the study, while continuing other measurements occurs frequently in practice since certain measurements may be costly or invasive. An example is the well-known Women's Health Initiative (WHI, Anderson et al., 2004; Rossouw et al., 2002), one of the largest multiple-endpoint randomized prevention studies of its kind. The WHI dropped the

endpoints designed to investigate the effect of hormone replacement therapy on cardiovascular and cancer outcomes in 2002 and 2005, respectively, but continued to follow-up participants for dementia and other cognition-related endpoints, known as the Women's Health Initiative Memory Study (Espeland et al., 2004; Shumaker et al., 1998). On the other hand, the setting of De and Baron (2012a,b) may be more appropriate when there is little motivation or cost-savings in dropping individual data streams, such as when the data (2) are the answers to k questions on a questionnaire filled out by the n th patient.

The novelty of the procedures of De and Baron (2012a,b) and those proposed herein is that they simultaneously control both the type I and II FWERs. For example, the procedure of Bartroff and Lai (2010) allows arbitrary (i.e., non-binding) acceptances of null hypotheses while controlling the type I FWER, hence the relationship between these acceptances and the power of the procedure is necessarily only available by analysis on a case-by-case basis. The current approach quantifies that relationship by, given a prescribed bound on the type II FWER, specifying an acceptance rule satisfying that bound. In some applications a prescribed value of the type II FWER may not be as readily motivated as the desired level of type I FWER. In this case we encourage the statistician to view the prescribed type II FWER as a parameter to choose in order to obtain a procedure with other desirable operating characteristics, such as average sample size in the streamwise or maximum sense.

2 General Formulation

2.1 Notation and Set-Up

For simplicity of presentation we introduce the procedure in the fully-sequential setting where the possible stopping times can be any positive integer $n = 1, 2, \dots$, although formulations in other settings like group-sequential and truncated settings are possible with only minor modifications. Fix the number $k \geq 2$ of data streams and let $\mathbf{k} = \{1, \dots, k\}$. Assume that there are k streams (1) of sequentially observable data and, for each $i \in \mathbf{k}$, it is desired to test the null hypothesis $H^{(i)}$ versus the alternative hypothesis $G^{(i)}$ about the parameter $\theta^{(i)}$ governing the i th data stream $X_1^{(i)}, X_2^{(i)}, \dots$, where $H^{(i)}$ and $G^{(i)}$ are disjoint subsets of the parameter space $\Theta^{(i)}$ containing $\theta^{(i)}$. The individual parameters $\theta^{(i)}$ may themselves be vectors, and the *global parameter* $\theta = (\theta^{(1)}, \dots, \theta^{(k)})$ is the concatenation of the individual parameters and is contained in the global parameter space $\Theta = \Theta^{(1)} \times \dots \times \Theta^{(k)}$.

Given $\theta \in \Theta$, let $\mathcal{T}(\theta) = \{i \in \mathbf{k} : \theta^{(i)} \in H^{(i)}\}$ denote the indices of the “true” hypotheses and $\mathcal{F}(\theta) = \{i \in \mathbf{k} : \theta^{(i)} \in G^{(i)}\}$ the indices of the “false” null hypotheses. The type I and II familywise error rates, denoted $\text{FWE}_I(\theta)$ and $\text{FWE}_{II}(\theta)$, are defined as

$$\begin{aligned}\text{FWE}_I(\theta) &= P_\theta(H^{(i)} \text{ is rejected for some } i \in \mathcal{T}(\theta)) \\ \text{FWE}_{II}(\theta) &= P_\theta(H^{(i)} \text{ is accepted for some } i \in \mathcal{F}(\theta)).\end{aligned}$$

Here the notion of rejecting (resp. accepting) $H^{(i)}$ is equivalent to accepting (resp. rejecting) $G^{(i)}$. This definition of $\text{FWE}_I(\theta)$ is the same as the standard one for fixed-sample testing (such as in Hochberg and Tamhane, 1987) and $\text{FWE}_{II}(\theta)$ is defined analogously; the

quantity $1 - \text{FWE}_{II}(\theta)$ has been called “familywise power” by some authors (e.g., Lee, 2004).

The building blocks of the sequential Holm procedure defined below are k individual sequential test statistics $\{\Lambda^{(i)}(n)\}_{i \in \mathbf{k}, n \geq 1}$, where $\Lambda^{(i)}(n)$ is the statistic for testing $H^{(i)}$ vs. $G^{(i)}$ based on the data $X_1^{(i)}, X_2^{(i)}, \dots, X_n^{(i)}$ available from the i th stream at time n . Concrete examples of these test statistics are given later in this section and in Section 3 but, for now, the reader may think of $\Lambda^{(i)}(n)$ as a sequential log likelihood ratio statistic for testing $H^{(i)}$ vs. $G^{(i)}$, for example. Given desired $\text{FWE}_I(\theta)$ and $\text{FWE}_{II}(\theta)$ bounds α and $\beta \in (0, 1)$, respectively, for each data stream i we assume the existence of critical values $A_s^{(i)} = A_s^{(i)}(\alpha, \beta)$ and $B_s^{(i)} = B_s^{(i)}(\alpha, \beta)$, $s \in \mathbf{k}$, such that

$$P_{\theta^{(i)}}(\Lambda^{(i)}(n) \geq B_s^{(i)} \text{ some } n, \Lambda^{(i)}(n') > A_1^{(i)} \text{ all } n' < n) \leq \frac{\alpha}{k-s+1} \quad \text{for all } \theta^{(i)} \in H^{(i)} \quad (3)$$

$$P_{\theta^{(i)}}(\Lambda^{(i)}(n) \geq A_s^{(i)} \text{ some } n, \Lambda^{(i)}(n') > B_1^{(i)} \text{ all } n' < n) \leq \frac{\beta}{k-s+1} \quad \text{for all } \theta^{(i)} \in G^{(i)} \quad (4)$$

for all $i, s \in \mathbf{k}$. We will show below that, in most cases, there are standard sequential statistics that satisfy these error bounds. Without loss of generality we assume that, for each $i \in \mathbf{k}$,

$$A_1^{(i)} \leq A_2^{(i)} \leq \dots \leq A_k^{(i)} < B_k^{(i)} \leq B_{k-1}^{(i)} \leq \dots \leq B_1^{(i)}. \quad (5)$$

For example, if the $A_s^{(i)}$ were not non-decreasing in s then they could be replaced by $\tilde{A}_s^{(i)} = \max\{A_1^{(i)}, \dots, A_s^{(i)}\}$ for which (4) would still hold; similarly for $B_s^{(i)}$ and (3). Note that, by (3)–(4), the critical values $A_1^{(i)}, B_1^{(i)}$ are simply the critical values for the sequential test that samples until

$$A_1^{(i)} < \Lambda^{(i)}(n) < B_1^{(i)} \quad (6)$$

is violated, and this test has type I and II error probabilities α/k and β/k , respectively. The values $A_s^{(i)}, s \in \mathbf{k}$, are then such that the similar sequential test with critical values $A_s^{(i)}, B_1^{(i)}$ has type II error probability $\beta/(k-s+1)$, and the analogous statement holds for the test with critical values $A_1^{(i)}, B_s^{(i)}$.

The sequential multiple testing procedure introduced below will involve ranking the test statistics associated with different data streams, which may be on completely different scales in general, so for each stream i we introduce a *standardizing function* $\phi^{(i)}(\cdot)$ which will be applied to the statistic $\Lambda^{(i)}(n)$ before ranking. The standardizing functions $\phi^{(i)}$ can be any increasing functions such that $\varphi^{(i)}(A_s^{(i)})$ and $\varphi^{(i)}(B_s^{(i)})$ do not depend on i . For simplicity, here we take the $\phi^{(i)}$ to be an increasing function satisfying

$$\varphi^{(i)}(A_s^{(i)}) = -(k-s+1) \quad \text{and} \quad \varphi^{(i)}(B_s^{(i)}) = k-s+1 \quad \text{for all } s \in \mathbf{k}. \quad (7)$$

For example, given $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$, an increasing piecewise linear function satisfying (7) can easily be constructed. Now let

$$\tilde{\Lambda}^{(i)}(n) = \varphi^{(i)}(\Lambda^{(i)}(n)) \quad (8)$$

and (3)–(4) can be written as

$$P_{\theta^{(i)}}(\tilde{\Lambda}^{(i)}(n) \geq s \text{ some } n, \tilde{\Lambda}^{(i)}(n') > -k \text{ all } n' < n) \leq \alpha/s \quad \text{for all } \theta^{(i)} \in H^{(i)} \quad (9)$$

$$P_{\theta^{(i)}}(\tilde{\Lambda}^{(i)}(n) \leq -s \text{ some } n, \tilde{\Lambda}^{(i)}(n') < k \text{ all } n' < n) \leq \beta/s \quad \text{for all } \theta^{(i)} \in G^{(i)} \quad (10)$$

for all $i, s \in \mathbf{k}$.

We shall describe the sequential Holm procedure in terms of *stages* of sampling, between which accept/reject decisions are made. Let $k_j \in \mathbf{k}$ ($j = 1, 2, \dots$) denote the number of *active* hypotheses (i.e., the $H^{(i)}$ which have been neither accepted nor rejected yet) at the beginning of the j th stage of sampling, and n_j will denote the cumulative sample size of any active test statistic up to and including the j th stage. The total number of null hypotheses that have been rejected (resp. accepted) so far at the beginning of the j th stage will be denoted by r_j (resp. a_j). Accordingly, set $k_1 = k$, $n_0 = 0$, $a_1 = r_1 = 0$, and fix desired $\text{FWE}_I(\theta)$ and $\text{FWE}_{II}(\theta)$ bounds α and β , respectively.

2.2 The Sequential Holm Procedure

The j th stage of sampling ($j = 1, 2, \dots$) proceeds as follows, in which we let $H^{(1)}, \dots, H^{(k_j)}$ denote the active hypotheses without loss of generality.

1. Sample the active streams $\{X_n^{(i)}\}_{i \in \mathbf{k}_j, n > n_{j-1}}$ until n equals

$$n_j = \inf \left\{ n > n_{j-1} : -(k - a_j) < \tilde{\Lambda}^{(i)}(n) < k - r_j \quad \text{is violated for some } i \in \mathbf{k}_j \right\}. \quad (11)$$

2. Relabel the active streams, if necessary, so that

$$\tilde{\Lambda}^{(1)}(n_j) \leq \tilde{\Lambda}^{(2)}(n_j) \leq \dots \leq \tilde{\Lambda}^{(k_j)}(n_j). \quad (12)$$

3. a. If the first inequality in (11) was violated, i.e., if $\tilde{\Lambda}^{(i)}(n_j) < -(k - a_j)$ for some $i \in \mathbf{k}_j$, then accept the $m_j - 1$ null hypotheses

$$H^{(1)}, H^{(2)}, \dots, H^{(m_j)},$$

where

$$m_j = \min \left\{ m \geq 1 : \tilde{\Lambda}^{(m+1)}(n_j) > -(k - a_j - m) \right\}, \quad (13)$$

and set $a_{j+1} = a_j + m_j$. Otherwise set $a_{j+1} = a_j$.

- b.** If the second inequality in (11) was violated, i.e., if $\tilde{\Lambda}^{(i)}(n_j) > k - r_j$ for some $i \in \mathbf{k}_j$, then reject the $m'_j \geq 1$ null hypotheses

$$H^{(k_j)}, H^{(k_j-1)}, \dots, H^{(k_j-m'_j+1)}, \quad (14)$$

where

$$m'_j = \min \left\{ m \geq 1 : \tilde{\Lambda}^{(k_j-m)}(n_j) < k - r_j - m \right\}, \quad (15)$$

and set $r_{j+1} = r_j + m'_j$. Otherwise set $r_{j+1} = r_j$.

- 4.** Stop if there are no remaining active hypotheses, i.e., if $a_{j+1} + r_{j+1} = k$. Otherwise, let k_{j+1} be the number of remaining active hypotheses and continue on to stage $j + 1$.

Before giving an example of this procedure, we make some remarks about its definition.

- A.** There will never be a conflict between the acceptances in Step 3a and the rejections in Step 3b since if $H^{(i)}$ is accepted at stage j and i is the index assigned in Step 2 of the j th stage, then $i \leq m_j$, hence $i - 1 < m_j$ so by (13) we have

$$\tilde{\Lambda}^{(i)}(n_j) \leq -(k - a_j - (i - 1)) < 0 < k - r_j - (k_j - i),$$

which shows that the set in (15) must contain the value $k_j - i$, hence $m'_j \leq k_j - i$, or $i \leq k_j - m'_j$. With (14), this shows that $H^{(i)}$ could not have also been rejected. A similar argument shows that a null hypothesis that is rejected could not also be accepted at the same stage.

- B.** If $k = 1$ then this definition becomes the sequential test (6) of the single null hypothesis $H^{(1)}$ versus alternative $G^{(1)}$ which has type I and II error probabilities bounded by α and β , respectively.
- C.** Ties in (12) can be broken arbitrarily (at random, say) without affecting the error control proved in Theorem 2.1, below.
- D.** If the same critical values are used for all data streams, that is, if $A_s^{(i)} = A_s^{(i')} = A_s$ and $B_s^{(i)} = B_s^{(i')} = B_s$ for all $i, i', s \in \mathbf{k}$, then the standardization performed in (8) can be dispensed with as long as the values compared with $\tilde{\Lambda}^{(i)}$ in the procedure's

definition are replaced by the appropriate values of A_s and B_s . Error control still holds under these conditions, which we prove below as part of Theorem 2.1.

- E.** The critical values $A_s^{(i)}, B_s^{(i)}$ can also depend on the sample size n of the test statistic being compared to them, with only notational changes in the definition of the procedure and the properties proved below. However, to avoid overly cumbersome notation we have omitted this from the presentation. Standard group sequential stopping boundaries – such as Pocock, O’Brien-Fleming, power family, and any others (see Jennison and Turnbull, 2000, Chapters 2 and 4) – can be utilized for the individual test statistics in this way.

Before stating our main result, Theorem 2.1, that this procedure controls both type I and II FWERs, we give a simplistic example to show the mechanics of the procedure. Table 1 contains three sample paths in the setting of three pairs of null and alternative hypotheses about the probability $p^{(i)}$ of success in Bernoulli data $X_n^{(i)}, i = 1, 2, 3$. Here the test statistics $\Lambda^{(i)}(n)$ are taken to be log likelihood ratios

$$\Lambda^{(i)}(n) = (2S_n^{(i)} - n) \log(.6/.4) \quad \text{where} \quad S_n^{(i)} = \sum_{j=1}^n X_j^{(i)}, \quad (16)$$

for testing

$$H^{(i)}: p^{(i)} \leq .4 \quad \text{vs.} \quad G^{(i)}: p^{(i)} \geq .6, \quad (17)$$

$i = 1, 2, 3$, about the success probability $p^{(i)}$ of i.i.d. Bernoulli data. This choice of test statistic and calculation of the critical values given in the table’s header will be explained in detail further below in Section 3.1.1; for now we merely focus on the procedure’s decisions to stop or continue sampling. Per remark (D) we dispense with the standardizing functions and drop the superscript (i) from the critical values A_s, B_s with which the statistics are compared, rather than the values $\mp(k - s + 1)$ from (7). The values of the stopped test statistics are given in bold in Table 1, and for clarity in this example we do not relabel the streams at each stage as in Step 2 of the procedure’s definition.

On sample path 1, sampling proceeds until time $n_1 = 7$ when $H^{(1)}$ and $H^{(2)}$ are rejected because this is the first time any of the 3 test statistics exceed B_1 or fall below A_1 . In particular, $H^{(1)}$ is rejected because $\Lambda^{(1)}(7) = 2.03$ $B_1 = 1.93$ and $H^{(2)}$ is also rejected at this time because $\Lambda^{(2)}(7) = 2.03$ $B_2 = 1.53$ and one null hypothesis (i.e., $H^{(1)}$) has already been rejected; the fact that $\Lambda^{(2)}(7)$ also exceeds B_1 was not necessary for rejecting $H^{(2)}$. Next, sampling of stream 3 is continued until time $n_2 = 10$ when $H^{(3)}$ is accepted because its test statistic falls below $A_1 = -2.43$. Similarly, on sample path 2, after rejecting $H^{(1)}$ at time $n_1 = 7$, $H^{(2)}$ is then rejected at time $n_2 = 8$ because $\Lambda^{(2)}(8)$ exceeds $B_2 = 1.53$ and one null hypothesis (i.e., $H^{(1)}$) has already been rejected. $H^{(3)}$ is also accepted at time $n_2 = 8$ for the same reason as above. On sample path 3, all three null hypotheses are rejected at time $n_1 = 7$ because $\Lambda^{(1)}(7) = 2.03$ B_1 , $\Lambda^{(2)}(7) = 2.03$ B_2 and one null hypothesis (i.e., $H^{(1)}$) has

already been rejected, and $\Lambda^{(3)}(7) = 1.22 - B_3$ and two null hypotheses (i.e., $H^{(1)}$ and $H^{(2)}$) have already been rejected.

Next we state the result that the sequential Holm procedure controls the type I and II FWERs, which is proved in the appendix.

Theorem 2.1—Fix $\alpha, \beta \in (0, 1)$. If the test statistics $\Lambda^{(i)}(n)$, $i \in \mathbf{k}$, $n \geq 1$, and critical values $A_s^{(i)} = A_s^{(i)}(\alpha, \beta)$ and $B_s^{(i)} = B_s^{(i)}(\alpha, \beta)$, $i, s \in \mathbf{k}$, satisfy (3)–(4), then the sequential Holm procedure defined above in Steps 1–4 satisfies $\text{FWE}_I(\theta) \leq \alpha$ and $\text{FWE}_{II}(\theta) \leq \beta$ for all $\theta \in \Theta$. If $A_s^{(i)} = A_s^{(i')} = A_s$ and $B_s^{(i)} = B_s^{(i')} = B_s$ for all $i, i' \in \mathbf{k}$, then this conclusion still holds if we take $\phi^{(i)}(x) = x$ for all $i \in \mathbf{k}$ and replace the values $\mp(k - s + 1)$ in the procedure's definition by A_s and B_s , respectively.

3 Constructing Test Statistics that Satisfy (3)–(4) for Individual Data Streams

Since all that is needed in the above construction of the sequential Holm procedure are sequential test statistics and critical values satisfying (3)–(4) for each data stream, in this section we show how to construct them in a few different settings and give some examples.

3.1 Simple Hypotheses and Their Use as Surrogates for Certain Composite Hypotheses

In this section we show how to construct the test statistics $\Lambda^{(i)}(n)$ and critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ satisfying (3)–(4) for any data stream i such that $H^{(i)}$ and $G^{(i)}$ are both simple hypotheses. This setting is of interest in practice because many more complicated composite hypotheses can be reduced to simple hypotheses. In this case the test statistics $\Lambda^{(i)}(n)$ will be taken to be log-likelihood ratios because of their strong optimality properties of the resulting sequential probability ratio test (SPRT); see Chernoff (1972). In order to express the likelihood ratio tests in simple form, we now make the additional assumption that each data stream $X_1^{(i)}, X_2^{(i)}, \dots$ constitutes independent and identically distributed data. However, we stress that this independence assumption is limited to *within* each stream so that, for example, elements of $X_1^{(i)}, X_2^{(i)}, \dots$ may be correlated with (or even identical to) elements of another stream $X_1^{(i')}, X_2^{(i')}, \dots$. We represent the simple null and alternative hypotheses $H^{(i)}$ and $G^{(i)}$ by the corresponding distinct density functions $h^{(i)}$ (null) and $g^{(i)}$ (alternative) with respect to some common σ -finite measure $\mu^{(i)}$. Formally, the parameter space $\Theta^{(i)}$ corresponding to this data stream is the set of all densities f with respect to $\mu^{(i)}$, and $H^{(i)}$ is considered *true* if the true density $f^{(i)}$ satisfies $f^{(i)} = h^{(i)} \mu^{(i)}$ -a.s., and is *false* if $f^{(i)} = g^{(i)} \mu^{(i)}$ -a.s. The SPRT for testing $H^{(i)} : f^{(i)} = h^{(i)}$ vs. $G^{(i)} : f^{(i)} = g^{(i)}$ with type I and II error probabilities α and β , respectively, utilizes the simple log-likelihood ratio test statistic

$$\Lambda^{(i)}(n) = \sum_{j=1}^n \log \left(\frac{g^{(i)}(X_j^{(i)})}{h^{(i)}(X_j^{(i)})} \right) \quad (18)$$

and samples sequentially until $\Lambda^{(i)}(n) \geq A(\alpha, \beta)$ or $\Lambda^{(i)}(n) \leq B(\alpha, \beta)$, where the critical values $A(\alpha, \beta)$ and $B(\alpha, \beta)$ satisfy

$$P_{h^{(i)}}(\Lambda^{(i)}(n) \geq B(\alpha, \beta) \text{ some } n, \Lambda^{(i)}(n') > A(\alpha, \beta) \text{ all } n' < n) \leq \alpha \quad (19)$$

$$P_{g^{(i)}}(\Lambda^{(i)}(n) \leq A(\alpha, \beta) \text{ some } n, \Lambda^{(i)}(n') < B(\alpha, \beta) \text{ all } n' < n) \leq \beta. \quad (20)$$

There are a few different options for computing $A(\alpha, \beta)$ and $B(\alpha, \beta)$ in practice. They may be computed numerically via Monte Carlo or normal approximation to the log-likelihood ratio (18), but the most widely-used method is to use the simple, closed-form *Wald-approximations*

$$A(\alpha, \beta) = \log\left(\frac{\beta}{1-\alpha}\right), \quad B(\alpha, \beta) = \log\left(\frac{1-\beta}{\alpha}\right). \quad (21)$$

See Hoel et al. (1971, Section 3.3.1) or Siegmund (1985) for a derivation. Although, in general, the inequalities in (19)–(20) only hold approximately when $A(\alpha, \beta)$ and $B(\alpha, \beta)$ are given by (21), Hoel et al. (1971) show that the actual type I and II error probabilities when using (21) can only exceed α or β by a negligibly small amount in the worst case, and the difference approaches 0 for small α and β , which is relevant in the present multiple testing situation where we will utilize Bonferroni-type cutdowns of α and β . In what follows in this section we adopt (21) and use these to construct the critical values $A_s^{(i)}, B_s^{(i)}$ of the sequential Holm procedure. The extensive simulations performed in Section 4 show that this does not lead to any exceedances of the desired FWERs. Alternative approaches would be to compute $\{A_s^{(i)}, B_s^{(i)}\}_{s \in k}$ via Monte Carlo, as mentioned above, or to replace (21) by $\log \beta$ and $\log \alpha^{-1}$, respectively, for which (19)–(20) always hold (see Hoel et al., 1971) and proceed similarly, but we do not explore those options here.

The next theorem shows that, neglecting Wald's approximation, the following simple expressions (22) can be used for the critical values in the sequential Holm procedure. Specifically, we show that the left-hand-sides of (3)–(4) are equal to the right-hand-sides of (23)–(24), and hence the inequalities in (3)–(4) hold, up to Wald's approximation.

Theorem 3.1: Suppose that, for a certain data stream i , $H^{(i)} : f^{(i)} = h^{(i)}$ and $G^{(i)} : f^{(i)} = g^{(i)}$ are simple hypotheses. Let $\alpha_{Wald}^{(i)}(\alpha, \beta)$ and $\beta_{Wald}^{(i)}(\alpha, \beta)$ be the values of the probabilities on the left-hand-sides of (19) and (20), respectively, with $\Lambda^{(i)}(n)$ given by (18) and $A(\alpha, \beta)$ and $B(\alpha, \beta)$ given by the Wald approximations (21). Now fix $\alpha, \beta \in (0, 1)$ and for $s \in k$ let

$$\alpha_s = \alpha_s(\alpha, \beta) = \frac{(k-s+1-\beta)\alpha}{(k-s+1)(k-\beta)}, \quad \beta_s = \beta_s(\alpha, \beta) = \frac{(k-s+1-\alpha)\beta}{(k-s+1)(k-\alpha)}.$$

Also, let $\alpha_{Holm}^{(i)}(s)$ and $\beta_{Holm}^{(i)}(s)$ denote the left-hand-sides of (3) and (4), respectively, with $A_s^{(i)}, B_s^{(i)}$ given by

$$A_s^{(i)} = A_s^{(i)}(\alpha, \beta) = \log \left(\frac{\beta}{(1-\alpha_s)(k-s+1)} \right), \quad B_s^{(i)} = B_s^{(i)}(\alpha, \beta) = \log \left(\frac{(1-\beta_s)(k-s+1)}{\alpha} \right). \quad (22)$$

Then, for all $s \in \mathbf{k}$,

$$\alpha_{Holm}^{(i)}(s) = \alpha_{Wald}^{(i)}(\alpha/(k-s+1), \beta_s) \quad \text{and} \quad (23)$$

$$\beta_{Holm}^{(i)}(s) = \beta_{Wald}^{(i)}(\alpha_s, \beta/(k-s+1)) \quad (24)$$

and therefore (3)–(4) hold, up to Wald's approximation, when using the critical values (22).

The theorem gives simple, closed form critical values (22) that can be used in lieu of Monte Carlo or other methods of calculating the $2k$ critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ for a stream i whose hypotheses $H^{(i)}, G^{(i)}$ are simple. Example values of (22) for $\alpha = .05$ and $\beta = .2$ are given in Table 2 for $k = 2, \dots, 10$.

3.1.1 Example: Exponential families—Suppose that a certain data stream i is comprised of i.i.d. d -dimensional random vectors $X_1^{(i)}, X_2^{(i)}, \dots$ from a multiparameter exponential family of densities

$$X_n^{(i)} \sim f_{\theta^{(i)}}(x) = \exp[\theta^{(i)T} x - \psi^{(i)}(\theta^{(i)})], \quad n=1, 2, \dots, \quad (25)$$

where $\theta^{(i)}$ and x are d -vectors, $(\cdot)^T$ denotes transpose, $\psi: \mathbb{R}^d \rightarrow \mathbb{R}$ is the cumulant generating function, and it is desired to test

$$H^{(i)}: \theta^{(i)} = \eta \quad \text{vs.} \quad G^{(i)}: \theta^{(i)} = \gamma$$

for given $\eta, \gamma \in \mathbb{R}^d$. Letting $S_n^{(i)} = \sum_{j=1}^n X_j^{(i)}$, the log-likelihood ratio (18) in this case is

$$\Lambda^{(i)}(n) = (\gamma - \eta)^T S_n^{(i)} - n[\psi^{(i)}(\gamma) - \psi^{(i)}(\eta)] \quad (26)$$

and, by Theorem 3.1, the critical values (22) can be used, which satisfy (3)–(4) up to Wald's approximation.

As mentioned above, many more complicated testing situations reduce to this setting. For example, the Bernoulli example (16) for testing (17) can be reduced to testing $p^{(i)} = .4$ vs. $p^{(i)} = .6$ by considering the worst-case error probabilities under the hypotheses (17), hence

(16) is given by (26) with $\theta^{(i)} = \log[p^{(i)}/(1 - p^{(i)})]$, $\psi^{(i)}(\theta^{(i)}) = -\log(1 - p^{(i)})$, $\eta = \log(.4/.6) = -\gamma$, and the critical values in Table 1 are given by (22) with $\alpha = \beta = .25$, this value chosen merely to produce short sample paths for the sake of the example.

3.2 Other Composite Hypotheses

While many composite hypotheses can be reduced to the simple-vs.-simple situation in Section 3.1, the generality of Theorem 2.1 does not require this and allows any type of hypotheses to be tested as long as the corresponding sequential statistics satisfy (3)–(4). In this section we discuss the more general case of how to proceed to apply Theorem 2.1 when a certain data stream i is described by a multiparameter exponential family (25) but simple hypotheses are not appropriate. Let

$$I(\theta^{(i)}, \lambda^{(i)}) = (\theta^{(i)} - \lambda^{(i)})^T \nabla \psi^{(i)}(\theta^{(i)}) - [\psi^{(i)}(\theta^{(i)}) - \psi^{(i)}(\lambda^{(i)})]$$

denote the Kullback-Leibler information number for the distribution (25), and it is desired to test

$$H^{(i)}: u(\theta^{(i)}) \leq u_0 \quad \text{vs.} \quad G^{(i)}: u(\theta^{(i)}) \geq u_1 \quad (27)$$

where $u(\cdot)$ is a continuously differentiable real-valued function such that

$$\text{for all fixed } \theta^{(i)}, I(\theta^{(i)}, \lambda^{(i)}) \text{ is } \begin{pmatrix} \text{decreasing} \\ \text{increasing} \end{pmatrix} \text{ in } u(\lambda^{(i)}) \begin{pmatrix} < \\ > \end{pmatrix} u(\theta^{(i)}),$$

and $u_0 < u_1$ are chosen real numbers. The family of models (25) and general form of the hypotheses (27) contain a large number of situations frequently encountered in practice, including various two-population tests that occur frequently in randomized Phase II and Phase III clinical trials and other situations involving nuisance parameters; see Bartroff et al. (2013, Chapter 4).

Of course there are many composite hypotheses encountered in practice which do not fit into the form (27), such as

$$\tilde{H}^{(i)}: \theta^{(i)} = \theta_0^{(i)} \quad \text{vs.} \quad \tilde{G}^{(i)}: \theta \neq \theta_0^{(i)} \quad (28)$$

for some fixed $\theta_0^{(i)}$. However, by considering true values of $\theta^{(i)}$ arbitrarily close to $\theta_0^{(i)}$, it is clear that no test of (28) can control the type II error probability for all $\theta^{(i)} \in \tilde{G}^{(i)}$ in general, and since the focus here is on tests that control both the type I and II FWERs, one would need to restrict $\tilde{G}^{(i)}$ in some way for that to be possible, for example by modifying $\tilde{G}^{(i)}$ to be only the $\theta^{(i)}$ such that $\|\theta^{(i)} - \theta_0^{(i)}\|^2 \geq \delta$ for some $\delta > 0$. But this restricted form fits into the

framework (27) by choosing $u(\theta^{(i)}) = \|\theta^{(i)} - \theta_0^{(i)}\|^2$, $u_0 = 0$, and $u_1 = \delta$. So although it is not natural to test (28) in the current framework of simultaneous type I and II FWER control, it is natural to test (28) when only type I FWER control is strictly required, and that problem has already been addressed in the sequential multiple testing setting by Bartroff and Lai (2010).

The hypotheses (27) can be tested with sequential generalized likelihood ratio (GLR) statistics, as follows. Letting

$$\hat{\theta}_n^{(i)} = (\nabla \psi^{(i)})^{-1} \left(\frac{1}{n} \sum_{j=1}^n X_j^{(i)} \right)$$

denote the maximum likelihood estimate (MLE) of θ based on the data from the first n observations, define

$$\Lambda_H(n) = n \left[\inf_{\lambda: u(\lambda) = u_0} I(\hat{\theta}_n^{(i)}, \lambda) \right], \quad (29)$$

$$\Lambda_G(n) = n \left[\inf_{\lambda: u(\lambda) = u_1} I(\hat{\theta}_n^{(i)}, \lambda) \right], \quad (30)$$

$$\Lambda^{(i)}(n) = \begin{cases} +\sqrt{2n\Lambda_H(n)}, & \text{if } u(\hat{\theta}_n^{(i)}) > u_0 \text{ and } \Lambda_H(n) \geq \Lambda_G(n) \\ -\sqrt{2n\Lambda_G(n)}, & \text{otherwise,} \end{cases} \quad (31)$$

where The statistics (29) and (30) are the log-GLR statistics for testing *against* $H^{(i)}$ and *against* $G^{(i)}$, respectively, whose signed roots in (31) have standard normal large- n limiting distribution under $u(\theta^{(i)}) = u_0$ and u_1 , respectively; see Jennison and Turnbull (1997, Theorem 2), whose results further show that under group sequential sampling, the signed-root statistics have asymptotically independent increments, a fact which can be used with random walk theory to find the critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ for $\Lambda^{(i)}(n)$ (see Bartroff et al., 2013, Chapter 4). However, our simulation studies have shown that under the fully-sequential sampling considered here, the small- n behavior of these statistics can deviate substantially from the standard normal random walk and therefore we advocate Monte Carlo determination of the critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ for $\Lambda^{(i)}(n)$, which then allows their inclusion in the sequential Holm procedure.

3.2.1 Example: Sequential Student's t -tests—As an example of the setup in this section for composite hypotheses, we consider sequential Student's t -tests. Suppose that the data $X_1^{(i)}, X_2^{(i)}, \dots$ from a certain data stream are i.i.d. Normal data with mean μ and variance σ^2 , both unknown, and it is desired to test the null hypothesis $\mu = 0$ versus the

alternative $\mu = \delta$ for some $\delta > 0$. Formally, this is a special case of the setup (27) by taking $\tilde{X}_j^{(i)} = (X_j^{(i)}, (X_j^{(i)})^2)^T$, $\theta^{(i)} = (\theta_1^{(i)}, \theta_2^{(i)})^T = (\mu/\sigma^2, -1/(2\sigma^2))^T$, $u(\theta^{(i)} = -\theta_1^{(i)}/(2\theta_2^{(i)})) = \mu$, $u_0 = 0$, and $u_1 = \delta$, and the parameter space and hypotheses are

$$\Theta^{(i)} = \mathbb{R} \times (-\infty, 0), \quad H^{(i)} = \left\{ (\theta_1, \theta_2)^T \in \Theta^{(i)} : \frac{-\theta_1}{2\theta_2} \leq 0 \right\}, \quad (32)$$

and $G^{(i)} = \left\{ (\theta_1, \theta_2)^T \in \Theta^{(i)} : \frac{-\theta_1}{2\theta_2} \geq \delta \right\}.$

By standard calculations in the exponential family (see Bartroff, 2006, p. 106), the log-GLR statistics (29)–(31) are

$$\Lambda_H(n) = (n/2) \log \left[1 + \left(\frac{\bar{X}_n^{(i)}}{\hat{\sigma}_n} \right)^2 \right], \quad (33)$$

$$\Lambda_G(n) = (n/2) \log \left[1 + \left(\frac{\bar{X}_n^{(i)} - \delta}{\hat{\sigma}_n} \right)^2 \right], \quad (34)$$

$$\Lambda^{(i)}(n) = \begin{cases} +\sqrt{2n\Lambda_H(n)}, & \text{if } \bar{X}_n^{(i)} \geq \delta/2 \\ -\sqrt{2n\Lambda_G(n)}, & \text{otherwise,} \end{cases} \quad (35)$$

where $\bar{X}_n^{(i)}$ and $\hat{\sigma}_n^2$ are the usual MLE estimates of μ and σ^2 , respectively, based on $X_1^{(i)}, \dots, X_n^{(i)}$.

The next lemma verifies that there are indeed critical values for which (3)–(4) hold, and that they can be found using only the standard Normal distribution. Using this result, recursive numerical integration can be used to find the critical values for which the error bounds hold conservatively, and this is the standard method used in this setting by the many sequential and group sequential software packages that exist; see Jennison and Turnbull (2000, Chapter 19) and Bartroff et al. (2013, Chapter 4.3).

Lemma 3.1: Let $A \leq 0 \leq B$ be arbitrary values and let $Z_1, Z_2, \dots$ be i.i.d. standard Normal

random variables with $T_n = \bar{Z}_n / \sqrt{\sum_{i=1}^n Z_i^2 / n}$, where $\bar{Z}_n$ and $\sum_{i=1}^n Z_i^2$ are the usual unbiased estimates of mean and variance, respectively, based on $Z_1, \dots, Z_n$. Then, in the sequential t -test setup above,

$$\sup_{\theta^{(i)} \in H^{(i)}} P_{\theta^{(i)}}(\Lambda^{(i)}(n) \geq B \text{ some } n, \Lambda^{(i)}(n') > A \text{ all } n' < n) \leq P(T_n \geq b_n \text{ some } n, T_{n'} > a_{n'} \text{ all } n' < n) \quad (36)$$

and

$$\sup_{\theta^{(i)} \in G^{(i)}} P_{\theta^{(i)}} (\Lambda^{(i)}(n) \leq A \text{ some } n, \Lambda^{(i)}(n') < B \text{ all } n' < n) \leq P(T_n \leq a_n \text{ some } n, T_{n'} < b_{n'} \text{ all } n' < n), \quad (37)$$

$$\text{where } a_n = -\sqrt{(n-1)(e^{(A/n)^2} - 1)} \text{ and } b_n = \sqrt{(n-1)(e^{(B/n)^2} - 1)}.$$

Therefore, for any fixed $\alpha, \beta \in (0, 1)$, there exist critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ such that (3)–(4) hold and which can be computed using only the standard Normal distribution of Z_j .

The lemma is proved in the Appendix.

In the sequential t -test setting, some authors (Lai and Zhang, 1994; Bartroff, 2006) have proposed bounding the standard deviation parameter σ from above by some finite value $\bar{\sigma} < \infty$. If this is the case then the parameter space and hypotheses (32) can be modified in the obvious way and it follows immediately from the proof of Lemma 3.1 in the Appendix that the derived boundaries a_n, b_n in (36)–(37) can be further refined by replacing them by $a_n + \delta_n$ and $\max\{b_n, \delta_n/2\}$, respectively, in (36) and by $\min\{a_n, -\delta_n/2\}$ and $b_n - \delta_n$, respectively, in (37), where $\delta_n = \delta \sqrt{n-1}/\bar{\sigma}$.

3.2.2 Example: Sequential Two-Population Binomial Tests—As another example of the setup in this section for composite hypotheses, we consider comparing two binomial populations. This setting is widely used in randomized phase II (Bartroff and Lai, 2008b, Section 3.2) and phase III (Bartroff and Lai, 2008a, Section 3.3) clinical trials and has recently been used to analyze RNA sequencing data (see Bartroff and Song, 2014).

Suppose that the data $X_1^{(i)}, X_2^{(i)}, \dots$ from a certain data stream is paired Binomial data $X_n^{(i)} = (Y_{1,n}^{(i)}, Y_{2,n}^{(i)})^T$ where the $Y_{j,n}^{(i)}$ ($j=1, 2$) are independent $\text{Bin}(m_j^{(i)}, p_j^{(i)})$ random variables, where the $m_j^{(i)}$ are the known “group” sizes and the $p_j^{(i)}$ are the unknown success probabilities. In some applications, particularly the RNA sequencing example mentioned above, the group sizes may vary sequentially (i.e., with n) and that situation can be handled with only minor modifications to what follows. Hypotheses that are a special case of (27) that can be tested are $H^{(i)}: p_1^{(i)} \leq p_2^{(i)}$ vs. $G^{(i)}: p_1^{(i)} \geq p_2^{(i)} + \delta$, for some $\delta > 0$. In particular,

$$\theta^{(i)} = \left(\log \left(\frac{p_1^{(i)}}{1-p_1^{(i)}} \right), \log \left(\frac{p_2^{(i)}}{1-p_2^{(i)}} \right) \right)^T, \\ u(\theta^{(i)}) = 1/(e^{-\theta_1^{(i)}} + 1) - 1/(e^{-\theta_2^{(i)}} + 1) = p_1^{(i)} - p_2^{(i)},$$

with $u_0 = 0$ and $u_1 = \delta$. The log GLR statistic (29) [resp. (30)] based on the data

$X_1^{(i)}, \dots, X_n^{(i)}$ takes the form

$$n \sum_{j=1}^2 \left[\hat{p}_j^{(i)} \log \left(\frac{\hat{p}_j^{(i)}}{\pi_j} \right) + (1 - \hat{p}_j^{(i)}) \log \left(\frac{1 - \hat{p}_j^{(i)}}{1 - \pi_j} \right) \right],$$

where $\hat{p}_j^{(i)} = \sum_{\ell=1}^n Y_{j,\ell} / (nm_j^{(i)})$ is the MLE of $p_j^{(i)}$ and π_j is the constrained MLE of $p_j^{(i)}$ subject to $p_1^{(i)} = p_2^{(i)}$ [resp. $p_1^{(i)} = p_2^{(i)} + \delta$].

4 Simulation Studies

In this section, we compare the sequential Holm procedure (denoted SH) with the fixed-sample Holm (1979) procedure (denoted FH), the sequential Bonferroni procedure (denoted SB), and the sequential intersection scheme (denoted IS) proposed by De and Baron (2012b). The SB procedure uses a SPRT on each data stream with error probability bounds α/k and β/k via the Wald approximations (21). That is, for each $i \in \mathbf{k}$, SB samples the i th stream until (6) is violated, with $A_1^{(i)} = \log[(\beta/k)/(1 - \alpha/k)]$ and $B_1^{(i)} = \log[(1 - \beta/k)/(\alpha/k)]$. The three sequential procedures SH, SB, and IS are the only ones we know of that control both FWE_I and FWE_{II} . In our studies we have chosen the commonly used values of $\alpha = .05$ and $\beta = .2$, i.e., familywise power at least 80%. This same value of α is used for the fixed-sample Holm procedure as well, which does not guarantee FWE_{II} control at a prescribed level, so we have chosen its sample size to make its familywise power approximately the same as that of the SH procedure in order to have a meaningful comparison with the sequential procedures. Below we present two sets of simulations, the first in Table 3 with independent streams of Bernoulli data, and the second in Table 4 with dependent streams of normal data generated from a multivariate normal distribution with non-identity covariance matrix. For each scenario considered below, FWE_I , FWE_{II} , expected total sample size

$EN = E(\sum_{i=1}^k N^{(i)})$ of all the data streams where $N^{(i)}$ is the total sample size of the i th stream, and relative savings in sample size of SH are estimated as the result of 100,000 Monte Carlo simulated batteries of k sequential tests. In each set of simulations, the data streams and hypotheses tested are similar for each data stream; we emphasize that this is only for the sake of getting a clear picture of the procedures' performance, and this uniformity is not required in order to be able to use the procedures considered.

We remark that we have included the IS procedure in the comparisons that follow only as a point of reference and, as mentioned in Section 1.1, it may be intended for situations in which no substantial cost savings is achieved by dropping individual data streams before sampling has been terminated. This should be kept in mind when viewing the expected total sample sizes of the various procedures below.

4.1 Setting 1: Independent Bernoulli Data Streams

Table 3 contains the operating characteristics of the above procedures for testing k hypotheses of the form

$$H^{(i)}:p^{(i)} \leq .4 \quad \text{vs.} \quad G^{(i)}:p^{(i)} \geq .6, \quad i=1, \dots, k,$$

about the probability $p^{(i)}$ of success in the i th stream of i.i.d. Bernoulli data; additionally, the streams were generated independently of each other. The individual test statistics (16) were used and the SH procedure used the critical values in Table 2, as described in Section 3.1. The data was generated for each data stream with $p^{(i)} = .4$ or $.6$ and the second column of Table 3 gives the number of hypotheses for which $p^{(i)} = .4$. The columns labeled Savings give the percent decrease in expected sample size EN of the SH relative to each other procedure. The SH procedure has substantially smaller sample size compared to the other three, saving more than 50% compared to FH and IS in each scenario with $k = 5$. Like its fixed-sample analog, the SB procedure is conservative in that its attained error rates FWE_I and FWE_{II} are much smaller than the prescribed levels, and the IS is similar in this regard, perhaps as a result of its stopping rule and resulting large average sample size. In fact, the IS has FWE_I and FWE_{II} even smaller than the SB procedure. The FH procedure has larger error rates than the SB and IS procedures, but still within the prescribed bounds, due to its step-down structure, and the SH procedure has very similar error rates to FH but with much smaller expected sample sizes.

4.2 Setting 2: Correlated Normal Data Streams

Table 4 contains the operating characteristics of the four procedures described above for testing k hypotheses of the form

$$H^{(i)}:\theta^{(i)} \leq 0 \quad \text{vs.} \quad G^{(i)}:\theta^{(i)} \geq \delta, \quad i=1, \dots, k,$$

for known $\delta > 0$, taken here to be 1, about the mean $\theta^{(i)}$ of i.i.d. normal observations with known variance 1, which makes up the i th data stream. To investigate the performance of the procedures under dependent data streams, the streams were generated from a k -dimensional multivariate normal distribution with mean $\theta = (\theta^{(1)}, \dots, \theta^{(k)})$, given in the third column of Table 4, and four different non-identity covariance matrices M_1, M_2, M_3 , and M_4 , given in the Appendix, which were chosen to give a variety of different scenarios of positively and negatively correlated data streams. The interaction of these various combinations of correlations with true or false null hypotheses all show somewhat similar behavior to the case of independent data streams in the previous section, in that the SH procedure has substantially smaller expected sample size than the other three procedures in all cases, more than a 30% reduction in most cases, and that the SH procedure has FWE_I and FWE_{II} much closer to the prescribed values α and β than the other two sequential procedures SB and IS, and similar to the FH procedure in most cases. Because the SH procedure causes more early stopping, it is interesting to note that its error control is less conservative than the other sequential procedures even in cases when data streams with true null hypotheses are positively correlated with streams having false null hypotheses, such as the second case of the M_1 -generated data and the third case of the M_3 -generated data

We have performed additional simulations in this setting but without the assumption of known variance, using the sequential t -test methodology in Section 3.2.1. The relative performance of the SH, FH, and SB procedures is similar to that seen in Table 4, with SH providing a large improvement over FH in terms of average sample size and a moderate improvement over SB, and all with slightly larger sample sizes due to the need to estimate the additional parameter, variance.

5 Discussion

The sequential Holm procedure proposed herein is a general method for combining individual sequential tests into a sequential multiple hypothesis testing procedure which controls both the type I and II FWERs at prescribed levels without requiring the statistician to have any knowledge or model of the data streams' correlation structure, a desirable property that it inherits from Holm's fixed-sample procedure. In our simulations in Section 4, the sequential Holm procedure exhibits much more efficiency in terms of smaller average total sample size than existing sequential procedures, as well as Holm's fixed-sample test. In terms of achieved FW-ERs, our simulations suggest that the sequential Holm procedure occupies a "middle ground" between existing sequential procedures, which have very conservative error rates and large average sample sizes, and the fixed-sample Holm test which achieves error rates closest to the prescribed values of all the procedures considered, but has still larger sample size and lacks the flexibility and adaptive nature of the sequential procedures.

We summarize our recommendations for using the sequential Holm procedure in practice as follows.

- For data streams whose hypotheses are simple, or are composite but can be reduced to considering simple hypotheses, we recommend using the sequential log-likelihood ratio statistic (18) with the closed-form critical values (22).
- For data streams with composite hypotheses of the form (27), we recommend using the sequential generalized likelihood ratio statistic (31) and determining the critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ to satisfy (3)–(4) by Monte Carlo. For group-sequential sampling with moderate group size the critical values can be determined by normal approximation.
- Data streams with still other forms of hypotheses or test statistics (e.g., nonparametric) can be included in the sequential Holm procedure by determining critical values $\{A_s^{(i)}, B_s^{(i)}\}_{s \in \mathbf{k}}$ satisfying (3)–(4) by Monte Carlo or other methods.

As mentioned in the introduction, this subject, which lays at the intersection of sequential analysis and multiple testing, is quite young and therefore still has many interesting and fundamental unanswered questions surrounding it. These include optimality theory for sequential multiple testing procedures, as well as calculations or estimates of their operating characteristics such as achieved FWERs and expected total and streamwise-maximum sample size.

Acknowledgments

The authors thank the Associate Editor and two reviewers for suggestions that improved the paper. The authors also thank Dr. Layla Parast for helpful discussions of the Women's Health Initiative.

Bartroff's work was partially supported by National Science Foundation grants DMS-0907241 and DMS-1310127, and National Institutes of Health grant GMS-068968. The majority of Song's work was completed while a PhD student in the Department of Mathematics at the University of Southern California.

References

- Anderson GL, Limacher MC, Assaf AR, Bassford T, Beresford SA, Black HR, Bonds DE, Brunner RL, Brzyski RG, Caan B, et al. Effects of conjugated equine estrogen in postmenopausal women with hysterectomy: The Women's Health Initiative randomized controlled trial. *Journal of the American Medical Association*. 2004; 291(14):1701–1712. [PubMed: 15082697]
- Baillie D. Multivariate acceptance sampling – some applications to defence procurement. *Journal of the Royal Statistical Society, Series D (The Statistician)*. 1987; 36(5):465–478.
- Bartroff, J. Recent Developments in Nonparametric Inference and Probability: Festschrift for Michael Woodroffe, volume 50 of IMS Lecture Notes Monograph Series. Hayward. Institute of Mathematical Statistics; 2006. Efficient three-stage t -tests; p. 105–111.
- Bartroff J, Lai TL. Efficient adaptive designs with mid-course sample size adjustment in clinical trials. *Statistics in Medicine*. 2008a; 27:1593–1611. [PubMed: 18275090]
- Bartroff J, Lai TL. Generalized likelihood ratio statistics and uncertainty adjustments in adaptive design of clinical trials. *Sequential Analysis*. 2008b; 27:254–276.
- Bartroff, J.; Lai, TL. Multistage tests of multiple hypotheses. In: Akahira, Honoring M.; Aoshima, M., editors. *Communications in Statistics – Theory and Methods*. Vol. 39. 2010. p. 1597–1607.
- Bartroff, J.; Lai, TL.; Shih, M. *Sequential Experimentation in Clinical Trials: Design and Analysis*. Springer; New York: 2013.
- Bartroff, J.; Song, J. Sequential tests of multiple hypotheses controlling false discovery and nondiscovery rates. 2014. Under review, <http://arxiv.org/abs/1311.3350>
- Chernoff, H. *Sequential Analysis and Optimal Design*. Society for Industrial and Applied Mathematics; Philadelphia: 1972.
- Cook R, Farewell V. Guidelines for monitoring efficacy and toxicity responses in clinical trials. *Biometrics*. 1994; 50:1146–1152. [PubMed: 7786995]
- De S, Baron M. Sequential Bonferroni methods for multiple hypothesis testing with strong control of family-wise error rates I and II. *Sequential Analysis*. 2012a; 31(2):238–262.
- De S, Baron M. Step-up and step-down methods for testing multiple hypotheses in sequential experiments. *Journal of Statistical Planning and Inference*. 2012b; 142:2059–2070.
- Dudoit, S.; van der Laan, M. *Multiple testing procedures with applications to genomics*. Springer Verlag; 2008.
- Espeland MA, Rapp SR, Shumaker SA, Brunner R, Manson JE, Sherwin BB, Hsia J, Margolis KL, Hogan PE, Wallace R, et al. Conjugated equine estrogens and global cognitive function in postmenopausal women: Womens Health Initiative Memory Study. *Journal of the American Medical Association*. 2004; 291(24):2959–2968. [PubMed: 15213207]
- Fisher, S. *Statistical Methods for Research Workers*. Oliver and Boyd; Edinburgh: 1932.
- Hochberg, Y.; Tamhane, AC. *Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics*. John Wiley & Sons Inc; New York: 1987. *Multiple Comparison Procedures*.
- Hoel, PG.; Port, SC.; Stone, CJ. *Introduction to Statistical Theory*. Houghton Mifflin Co; Boston, Mass: 1971.
- Holm S. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*. 1979; 6:65–70.
- Jennison C, Turnbull B. Group sequential tests for bivariate response: Interim analyses of clinical trials with both efficacy and safety endpoints. *Biometrics*. 1993; 49:741–752. [PubMed: 8241370]

- Jennison C, Turnbull BW. Group sequential analysis incorporating covariate information. *Journal of the American Statistical Association*. 1997; 92:1330–1341.
- Jennison, C.; Turnbull, BW. *Group Sequential Methods with Applications to Clinical Trials*. Chapman & Hall/CRC; New York: 2000.
- Lai TL, Zhang L. A modification of Schwarz's sequential likelihood ratio tests in multivariate sequential analysis. *Sequential Analysis*. 1994; 13:79–96.
- Lee, M. *Analysis of Microarray Gene Expression Data*. Springer; 2004.
- Lehmann, EL.; Romano, JP. *Springer Texts in Statistics*. 3. Springer; New York: 2005. *Testing Statistical Hypotheses*.
- Mei Y. Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*. 2010; 97(2):419–433.
- O'Brien PC. Procedures for comparing samples with multiple endpoints. *Biometrics*. 1984; 40:1079–1087. [PubMed: 6534410]
- O'Brien PC, Fleming TR. A multiple testing procedure for clinical trials. *Biometrics*. 1979; 35:549–556. [PubMed: 497341]
- Paulson E. A sequential procedure for selecting the population with the largest mean from k normal populations. *The Annals of Mathematical Statistics*. 1964; 35:174–180.
- Romano JP, Wolf M. Stepwise multiple testing as formalized data snooping. *Econometrica*. 2005; 73(4):1237–1282.
- Rossouw JE, Anderson GL, Prentice RL, LaCroix AZ, Kooperberg C, Stefanick ML, Jackson RD, et al. Risks and benefits of estrogen plus progestin in healthy postmenopausal women: principal results from the Women's Health Initiative randomized controlled trial. *Journal of the American Medical Association*. 2002; 288(3):321–333. [PubMed: 12117397]
- Scheffé H. A method for judging all contrasts in the analysis of variance. *Biometrika*. 1953; 40:87–110.
- Seber, GAF.; Lee, AJ. *Wiley Series in Probability and Statistics*. 2. Wiley-Interscience [John Wiley & Sons]; Hoboken, NJ: 2003. *Linear Regression Analysis*.
- Shumaker SA, Reboussin BA, Espeland MA, Rapp SR, McBee WL, Dailey M, Bowen D, Terrell T, Jones BN. The Womens Health Initiative Memory Study (WHIMS): A trial of the effect of estrogen therapy in preventing and slowing the progression of dementia. *Controlled Clinical Trials*. 1998; 19(6):604–621. [PubMed: 9875839]
- Siegmund, D. *Sequential Analysis: Tests and Confidence Intervals*. Springer-Verlag; New York: 1985.
- Siegmund D. A sequential clinical trial for comparing three treatments. *The Annals of Statistics*. 1993; 21(1):464–483.
- Tang DI, Geller NL, Pocock SJ. On the design and analysis of randomized clinical trials with multiple endpoints. *Biometrics*. 1993; 49:23–30. [PubMed: 8513104]
- Tang DI, Gnecco C, Geller NL. Design of group sequential clinical trials with multiple endpoints. *Journal of the American Statistical Association*. 1989; 84:776–779.
- Tartakovsky A, Li X, Yaralov G. Sequential detection of targets in multichannel systems. *Information Theory, IEEE Transactions on*. 2003; 49(2):425–445.
- Wald, A. *Sequential Analysis*. Wiley; New York: 1947. Reprinted by Dover, 1973
- Ye Y, Li A, Liu L, Yao B. A group sequential Holm procedure with multiple primary endpoints. *Statistics in Medicine*. 2013; 32:1112–1124. [PubMed: 23239078]

Appendix: Proofs and Details of Simulation Studies

Proofs

Proof of Theorem 2.1

We fix θ and, for simplicity, omit it from the notation that follows. We first prove that $\text{FWE}_{II} \leq \beta$. If $\mathcal{F} = \emptyset$ then $\text{FWE}_{II} = 0$, so assume that $\mathcal{F} \neq \emptyset$. Let

$$\mathcal{S}_j = \left\{ i \in \mathbf{k}_j : \tilde{\Lambda}^{(i)}(n_j) \leq -(k - a_j - i + 1) \text{ and } H^{(i)} \text{ false} \right\}, \quad (38)$$

where the indexing of the $\tilde{\Lambda}^{(i)}$ and $H^{(i)}$ in (38) is with respect to the ordering in Step 2 of the j th stage. Let j^* denote the earliest stage at which a false hypothesis is accepted, taking the value ∞ if no such error occurs; to prove that $\text{FWE}_{II} \leq \beta$ we thus assume without loss of generality that $j^* < \infty$ with probability 1. By our assumptions and by definition of j^* , $\mathcal{S}_{j^*} \neq \emptyset$ so let $i^* = \min \mathcal{S}_{j^*}$. Let ϕ_{j^*} be the number of false hypotheses active at the beginning of stage j^* , and $\phi_j^* = |\mathcal{F}| - \phi_j$ the number of false hypotheses rejected at some stage prior to j^* , where $|\cdot|$ denotes set cardinality. Clearly $\phi_{j^*} \leq r_{j^*}$, and by definition of i^* and j^* ,

$$\phi_{j^*} \leq k_{j^*} - (i^* - 1) = (k - a_{j^*} - r_{j^*}) - (i^* - 1).$$

Combining these gives

$$|\mathcal{F}| = \phi_{j^*} + \phi_{<j^*} \leq k - a_{j^*} - i^* + 1. \quad (39)$$

For the remainder of the proof, let the generic index $i \in \mathbf{k}$ denote the fixed, original indexing of data streams rather than an ordered index assigned in Step 2 of the procedure, but let all other indices be ordered. Let V_i be the event that $H^{(i)}$ is accepted at stage j^* and let i^* denote the index assigned to $\tilde{\Lambda}^{(i)}(n)$ in Step 2 of the j^* th stage. Then

$$V_i \subseteq \left\{ \tilde{\Lambda}^{(i)}(n_{j^*}) \leq -(k - a_{j^*} - i^* + 1) \right\} \subseteq \left\{ \tilde{\Lambda}^{(i)}(n_{j^*}) \leq -|\mathcal{F}| \right\} \quad (40)$$

by (39). We will also show that

$$V_i \subseteq \left\{ \tilde{\Lambda}^{(i)}(n) < k \text{ for all } n < n_{j^*} \right\}. \quad (41)$$

This holds because if $\tilde{\Lambda}^{(i)}(n) \geq k$ for some $n < n_{j^*}$, then $H^{(i)}$ would be rejected at some stage prior to j^* . To see this, let W_i be the event on the right-hand-side of (41). It is clear from Step 3b that on W_i^c , some hypothesis would be rejected at a stage $j < j^*$ since $k \leq k - r_j$ for any value of r_j , and this is the rejection boundary for statistics in (11). Let $j' < j^*$ denote the earliest stage such that $H^{(i)}$ is not rejected before stage j' and

$$\tilde{\Lambda}^{(i)}(n_{j'}) \geq k. \quad (42)$$

Let $k_{j'} - i'$ denote the ordered index assigned to $\tilde{\Lambda}^{(i)}(n_{j'})$ in Step 2 of the j' th stage. We will show that

$$\tilde{\Lambda}^{(k_{j'}-m)}(n_{j'}) \geq k-r_{j'}-m \quad \text{for all } 1 \leq m \leq i'$$

which, by (15), implies that $H^{(i)}$ is rejected at stage j' and finishes the proof of (41). For any $1 \leq m \leq i'$,

$$\tilde{\Lambda}^{(k_{j'}-m)}(n_{j'}) \geq \tilde{\Lambda}^{(k_{j'}-i')}(n_{j'}) \geq k \geq k-r_{j'}-m,$$

by (42).

Combining (40) and (41) we have

$$V_i \subseteq \left\{ \tilde{\Lambda}^{(i)}(n) \leq -|\mathcal{F}| \text{ some } n, \tilde{\Lambda}^{(i)}(n') < k \text{ all } n' < n \right\},$$

and thus $P(V_i) \leq \beta/|\mathcal{F}|$ by (10). Using this we have

$$\text{FWE}_{II} = P\left(\bigcup_{i \in \mathcal{F}} V_i\right) \leq \sum_{i \in \mathcal{F}} P(V_i) \leq \sum_{i \in \mathcal{F}} \beta/|\mathcal{F}| = \beta.$$

The proof that $\text{FWE}_I = \alpha$ is similar so the details are omitted. The only thing that could make the situation different is the possibility that a hypothesis that would have been rejected in Step 3b is accepted in Step 3a. However, Remark (A) guarantees that this does not happen.

For the second claim of the theorem, since $\phi^{(i)}(\cdot)$ is strictly increasing, its inverse could be applied to $\tilde{\Lambda}^{(i)}$ and the corresponding critical value at every place they are compared in the procedure's definition and the proof, thus replacing $\tilde{\Lambda}^{(i)}$ by $\Lambda^{(i)}$ and $\mp(k-s+1)$ by

$A_s^{(i)}, B_s^{(i)}$. If $A_s^{(i)} = A_s^{(i')} = A_s$ and $B_s^{(i)} = B_s^{(i')} = B_s$ for all $i, i', s \in \mathbf{k}$, then this is equivalent to using $\phi^{(i)}(x) = x$ for all $i \in \mathbf{k}$.

Proof of Theorem 3.1

First note that $\alpha_s, \beta_s \in (0, 1)$ for all $s \in \mathbf{k}$ since

$$0 < \alpha_s = \frac{k-s+1-\beta}{k-s+1} \cdot \frac{\alpha}{k-\beta} < 1 \cdot \frac{\alpha}{k-\beta} < \frac{1}{k-1} \leq 1$$

as $k \geq 2$, and similarly for β_s . $A_s^{(i)}$ and $B_s^{(i)}$ in (22) can be written as $A(\alpha_s, \beta'(k-s+1))$ and $B(\alpha/(k-s+1), \beta_s)$, respectively, and it is simple algebra to then check that

$A(\alpha/(k-s+1), \beta_s) = A_1^{(i)}$ for any $s \in \mathbf{k}$. Then, to verify (23),

$$\begin{aligned} \alpha_{Holm}^{(i)}(s) &= P_{h^{(i)}}(\Lambda^{(i)}(n) \geq B_s^{(i)} \text{ some } n, \Lambda^{(i)}(n') > A_1^{(i)} \text{ all } n' < n) \\ &= P_{h^{(i)}}(\Lambda^{(i)}(n) \geq B(\alpha/(k-s+1), \beta_s) \text{ some } n, \Lambda^{(i)}(n') > A(\alpha/(k-s+1), \beta_s) \text{ all } n' < n) \\ &= \alpha_{Wald}^{(i)}(\alpha/(k-s+1), \beta_s) \end{aligned}$$

by (19). The proof of (24) is similar.

Proof of Lemma 3.1

We will prove (36); the proof of (37) is similar after replacing $X_j^{(i)}$ by $\delta - X_j^{(i)}$. Let

$\tilde{\sigma}_n^2 = n\hat{\sigma}_n^2/(n-1)$ and $\tilde{T}_n = \sqrt{n\tilde{X}_n^{(i)}}/\tilde{\sigma}_n$. By (33)–(35),

$$\{\Lambda^{(i)}(n) \geq B\} = \{\bar{X}_n^{(i)} \geq \delta/2 \quad \text{and} \quad \sqrt{2n\Delta_H(n)} \geq B\} \subseteq \left\{ \frac{\bar{X}_n^{(i)}}{\hat{\sigma}_n} \geq \frac{b_n}{\sqrt{n-1}} \right\} = \{\tilde{T}_n \geq b_n\}$$

and

$$\begin{aligned} \{\Lambda^{(i)}(n) > A\} &= \{\bar{X}_n^{(i)} \geq \delta/2\} \cup \{\bar{X}_n^{(i)} < \delta/2 \quad \text{and} \quad -\sqrt{2n\Lambda_G(n)} > A\} \\ &= \{\bar{X}_n^{(i)} \geq \delta/2\} \cup \{\delta + \hat{\sigma}_n \frac{a_n}{\sqrt{n-1}} < \bar{X}_n^{(i)} < \delta/2\} \subseteq \left\{ \frac{\bar{X}_n^{(i)}}{\hat{\sigma}_n} > \frac{a_n}{\sqrt{n-1}} \right\} = \{\tilde{T}_n > a_n\}. \end{aligned}$$

Then for any $\theta^{(i)} = (\theta_1^{(i)}, \theta_2^{(i)})^T \in H^{(i)}$, the probability on the left-hand-side of (36) is bounded above by

$$P_{(0, \theta_2^{(i)})^T}(\tilde{T}_n \geq b_n \text{ some } n, \tilde{T}_{n'} > a_{n'} \text{ all } n' < n) \quad (43)$$

and since the distribution of T_n under $\theta^{(i)} = (0, \theta_2^{(i)})^T$ does not depend on $\theta_2^{(i)}$, we may take $\theta_2^{(i)} = -1/2$, i.e., the standard normal distribution, and thus (43) is equal to the right-hand-side of (36), completing the proof.

Details of Simulation Studies

The four covariance matrices used in the simulations for Section 4.2 are as the following:

$$\begin{aligned}
 M_1 &= \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix} \\
 M_2 &= \begin{pmatrix} 1 & -0.8 \\ -0.8 & 1 \end{pmatrix} \\
 M_3 &= \begin{pmatrix} 1 & 0.8 & -0.6 & -0.8 \\ 0.8 & 1 & -0.6 & -0.8 \\ -0.6 & -0.6 & 1 & 0.8 \\ -0.8 & -0.8 & 0.8 & 1 \end{pmatrix} \\
 M_4 &= \begin{pmatrix} 1 & 0.8 & 0.6 & -0.4 & -0.6 & -0.8 \\ 0.8 & 1 & 0.8 & -0.4 & -0.6 & -0.8 \\ 0.6 & 0.8 & 1 & -0.4 & -0.6 & -0.8 \\ -0.4 & -0.4 & -0.4 & 1 & 0.8 & 0.6 \\ -0.6 & -0.6 & -0.6 & 0.8 & 1 & 0.8 \\ -0.8 & -0.8 & -0.8 & 0.6 & 0.8 & 1 \end{pmatrix}
 \end{aligned}$$

Table 1

Three sample paths of the sequential Holm procedure for $k = 3$ hypotheses about Bernoulli data using critical values $A_1 = -2.34$, $A_2 = -1.94$, $A_3 = -1.27$, $B_1 = 1.93$, $B_2 = 1.53$, $B_3 = .86$. The values of the stopped sequential statistics are in bold.

Data Stream	$n = 1$	2	3	4	5	6	7	8	9	10
Sample Path 1										
1	$X_n^{(1)}$	0	1	1	1	1	1	1		
	$\Lambda^{(1)(n)}$	-.41	.00	.41	.81	1.22	1.62	2.03		
2	$X_n^{(2)}$	1	0	1	1	1	1	1		
	$\Lambda^{(2)(n)}$	.41	.00	.41	.81	1.22	1.62	2.03		
3	$X_n^{(3)}$	0	1	0	0	1	0	0	0	0
	$\Lambda^{(3)(n)}$	-.41	.00	-.41	-.81	-.41	-.81	-1.22	-1.62	-2.03
Sample Path 2										
1	$X_n^{(1)}$	0	1	1	1	1	1	1		
	$\Lambda^{(1)(n)}$	-.41	.00	.41	.81	1.22	1.62	2.03		
2	$X_n^{(2)}$	1	0	0	1	1	1	1	1	
	$\Lambda^{(2)(n)}$	.41	.00	-.41	.00	.41	.81	1.22	1.62	
3	$X_n^{(3)}$	0	1	0	0	0	0	0	0	
	$\Lambda^{(3)(n)}$	-.41	.00	-.41	-.81	-1.22	-1.62	-2.03	-2.43	
Sample Path 3										
1	$X_n^{(1)}$	1	0	1	1	1	1	1		
	$\Lambda^{(1)(n)}$	.41	.00	.41	.81	1.22	1.62	2.03		
2	$X_n^{(2)}$	1	1	1	0	1	1	1		
	$\Lambda^{(2)(n)}$	.41	.81	1.22	.81	1.22	1.62	2.03		
3	$X_n^{(3)}$	0	1	0	1	1	1	1		
	$\Lambda^{(3)(n)}$	-.41	.00	-.41	.00	.41	.81	1.22		

Table 2

Critical values (22) of the sequential Holm procedure for simple hypotheses, for $\alpha = .05$, $\beta = .2$, and $k = 2, \dots, 10$ to two decimal places.

k	$A_1, \dots, A_k$ $B_1, \dots, B_k$			
2	-2.28	-1.59		
	3.58	2.89		
3	-2.69	-2.29	-1.60	
	4.03	3.62	2.93	
4	-2.98	-2.70	-2.29	-1.60
	4.33	4.04	3.64	2.95
5	-3.21	-2.99	-2.70	-2.29
	4.56	4.34	4.05	3.65
				2.96
6	-3.39	-3.21	-2.99	-2.70
	4.75	4.57	4.35	4.06
				3.66
				2.96
7	-3.55	-3.39	-3.21	-2.99
	4.91	4.76	4.58	4.35
				4.07
				3.66
				2.97
8	-3.68	-3.55	-3.39	-3.21
	5.05	4.92	4.76	4.58
				4.36
				4.07
				3.66
				2.97
9	-3.80	-3.68	-3.55	-3.40
	5.17	5.05	4.92	4.77
				4.58
				4.36
				4.07
				3.67
				2.97
10	-3.91	-3.80	-3.68	-3.55
	5.28	5.17	5.05	4.92
				4.77
				4.59
				4.36
				4.07
				3.67
				2.98

Table 3

Operating characteristics of sequential and fixed-sample multiple testing procedures for k streams of independent Bernoulli data.

k	# of true $H^{(i)}$	SH				FH				SB				IS			
		FWE_I	FWE_{II}	EN	FWE_I	FWE_{II}	EN	Savings	FWE_I	FWE_{II}	EN	Savings	FWE_I	FWE_{II}	EN	Savings	
1	1	0.048	-	17.5	-	-	-	-	0.048	-	17.5	0.0%	0.031	-	18.1	3.8%	
	0	-	0.190	24.6	-	0.194	42	41.5%	-	0.190	24.6	0.0%	-	0.194	28.6	14.1%	
2	2	0.045	-	47.6	-	-	-	-	0.048	-	56.3	15.4%	0.025	-	64.7	26.4%	
	1	0.029	0.135	63.0	0.029	0.135	126	50.0%	0.025	0.086	66.7	5.6%	0.021	0.120	97.2	35.2%	
	0	-	0.165	72.7	-	0.168	126	42.3%	-	0.161	77.1	5.8%	-	0.091	104.0	30.1%	
5	3	0.034	0.105	216.7	0.039	0.108	485	55.3%	0.022	0.077	230.2	5.9%	0.010	0.044	474.9	54.4%	
	2	0.028	0.127	230.7	0.033	0.128	490	52.9%	0.015	0.112	247.1	6.6%	0.012	0.044	439.9	47.6%	
10	8	0.034	0.070	479.9	0.029	0.075	1200	60.0%	0.027	0.034	532.6	9.9%	0.011	0.026	1144.8	58.1%	
	5	0.027	0.111	549.6	0.045	0.112	1240	55.7%	0.017	0.085	587.1	6.4%	0.008	0.023	1295.0	57.6%	
	2	0.016	0.130	579.4	0.033	0.132	1180	50.9%	0.007	0.129	642.8	9.9%	0.006	0.022	1298.2	55.4%	
20	16	0.035	0.067	1129.8	0.047	0.072	2860	60.5%	0.035	0.029	1250.8	9.7%	0.006	0.011	3095.6	63.5%	
	10	0.027	0.108	1273.2	0.045	0.113	3040	58.1%	0.022	0.073	1336.5	4.7%	0.004	0.013	3406.0	62.6%	
	4	0.017	0.137	1332.6	0.035	0.138	2740	51.4%	0.009	0.116	1421.9	6.3%	0.004	0.015	3344.1	60.2%	

Table 4
Operating characteristics of sequential and fixed-sample multiple testing procedures for k streams of correlated Normal data.

			SH			FH			SB			IS					
Covariance	k	true θ	FWE $_I$	FWE $_{II}$	EN	FWE $_I$	FWE $_{II}$	EN	Savings	FWE $_I$	FWE $_{II}$	EN	Savings	FWE $_I$	FWE $_{II}$	EN	Savings
M_1	2	(1, 1)	0.024	-	10.4	-	-	-	-	0.025	-	11.6	10.0%	0.009	-	11.7	11.3%
		(1, 0)	0.029	0.110	12.8	0.050	0.113	20	35.9%	0.015	0.057	13.6	5.7%	0.027	0.108	19.5	34.3%
		(0, 0)	-	0.087	14.3	-	0.102	20	28.5%	-	0.086	15.6	8.5%	-	0.046	16.7	14.4%
M_2	2	(1, 1)	0.029	-	10.2	-	-	-	-	0.030	-	11.6	11.9%	0.025	-	15.1	32.5%
		(1, 0)	0.015	0.063	13.5	0.037	0.082	22	38.8%	0.015	0.057	13.6	0.8%	0.009	0.016	17.9	24.8%
		(0, 0)	-	0.114	14.4	-	0.128	20	28.1%	-	0.113	15.6	8.0%	-	0.103	20.3	29.3%
M_3	4	(1, 1, 1, 1)	0.024	-	24.6	-	-	-	-	0.025	-	29.1	15.6%	0.007	-	38.8	36.7%
		(1, 1, 0, 0)	0.013	0.051	32.4	0.032	0.058	60	46.0%	0.0011	0.044	33.8	4.0%	0.001	0.003	46.9	30.9%
		(1, 0, 1, 0)	0.020	0.080	32.2	0.043	0.102	56	42.5%	0.015	0.058	33.8	4.7%	0.010	0.021	59.6	46.0%
M_4	6	(0, 0, 0, 0)	-	0.089	34.1	-	0.115	48	29.1%	-	0.090	38.4	11.4%	-	0.029	52.6	35.3%
		(1, 1, 1, 1, 1, 1)	0.022	-	40.7	-	-	-	-	0.022	-	48.8	16.5%	0.002	-	67.1	39.2%
		(1, 1, 1, 1, 1, 0)	0.021	0.032	46.3	0.038	0.032	108	57.2%	0.020	0.019	51.2	9.6%	0.003	0.002	85.2	45.7%
M_4	6	(1, 1, 1, 1, 0, 0)	0.018	0.041	50.4	0.038	0.041	108	53.3%	0.016	0.031	53.7	6.1%	0.007	0.001	90.0	44.0%
		(1, 1, 0, 1, 1, 0)	0.018	0.068	53.4	0.039	0.073	102	47.7%	0.013	0.049	56.1	4.9%	0.004	0.014	102.3	47.8%
		(1, 1, 1, 0, 0, 0)	0.012	0.047	53.6	0.030	0.058	102	47.5%	0.012	0.041	56.2	4.7%	<0.001	<0.001	83.0	35.4%
M_4	6	(1, 1, 0, 0, 0, 0)	0.011	0.072	55.5	0.031	0.089	96	42.2%	0.008	0.061	58.6	5.3%	<0.001	0.022	97.8	43.3%
		(1, 0, 0, 1, 0, 0)	0.016	0.074	55.3	0.045	0.092	96	42.4%	0.010	0.059	58.7	5.7%	0.005	0.008	104.1	46.9%
		(1, 0, 0, 0, 0, 0)	0.008	0.077	56.3	0.040	0.099	90	37.5%	0.005	0.071	61.1	7.9%	0.001	0.011	97.8	42.4%
M_4	6	(0, 0, 0, 0, 0, 0)	-	0.082	55.7	-	0.087	84	33.7%	-	0.082	63.6	12.5%	-	0.009	89.4	37.8%