1 Sequential Generalized Likelihood Ratio Tests for Online Item Monitoring

2 Hyeon-Ah Kang

3 University of Texas at Austin

Abstract

The study presents statistical procedures that monitor functioning of items over time. We propose generalized likelihood ratio tests that surveil multiple item parameters and implement with various sampling techniques to perform continuous or intermittent monitoring. The procedures examine stability of item parameters across time and inform compromise as soon as they identify significant parameter shift. The performance of the monitoring procedures was validated using simulated and real-assessment data. The empirical evaluation suggests that the proposed procedures perform adequately well in identifying the parameter drift. They showed satisfactory detection power and gave timely signals while regulating the error rates reasonably low. The procedures also showed superior performance when compared with the existent methods. The empirical findings suggest that multivariate parametric monitoring can provide an efficient and powerful control tool for maintaining the quality of items. The procedures allow joint monitoring of multiple item parameters and achieve sufficient power by dint of likelihood-ratio tests. Based on the findings from the empirical experimentation, we suggest some practical strategies for performing online item monitoring.

Keywords: item parameter drift, online monitoring, sequential generalized likelihood ratio test, cumulative sum control chart, response time

Sequential Generalized Likelihood Ratio Tests for Online Item Monitoring

1. Introduction

Many testing programs use items iteratively for multiple test forms or continuous testing. When an item is used repeatedly, it can sometimes exhibit differential performance, displaying fluctuations in the item parameters. A systematic change in the item parameters—a phenomenon called item parameter drift (Goldstein, 1983)—can adversely affect validity and comparability of test scores (Huggins-Manley, 2017; Wells, Subkoviak, & Serlin, 2002). To minimize adverse ramifications on the testing outcomes, any significant parameter change must be identified as soon as it develops.

While the root cause of parameter drift is difficult to ascertain, studies have examined some probable reasons that can be offered to the drifts in real assessments (see Clark (2013) for review). One of the commonly cited reasons has been security breach (e.g., item-sharing, test collusion, fraud, cheating). For example, when an item is leaked to prospective examinees as a result of earlier test-takers' item-sharing, beneficiaries of the item preknowledge tend to exhibit distinct behavioral patterns, such as increased chance of item success and shorter response time. The consistent change in the examinee behaviors can alter item's operational properties and lead to systematic shift in the item parameters. Other reasons have also been contemplated to account for parameter drift in real tests. At other times, parameter drift has been attributed to change in the test's properties, such as evolution of constructs or content areas, and adjustments made to items (e.g., format, presentation mode, position). Temporal factors, such as seasonality and demographic variation, have also been suggested as plausible sources of drift.

While different reasons have been explored to account for drift in the item parameters, detection of drift has been achieved by similar methods. One salient approach has been recurrent item calibration. It re-estimates item parameters several times and compares new parameter estimates with the known values (e.g., Donoghue & Isham, 1998; Veerkamp & Glas, 2000). If the new estimates are significantly different from the known

values, it is claimed that the item parameters have drifted. In other studies, fluctuation in the item parameters has been directly modeled within the item response model (e.g., Bock, Muraki, & Pfeifferberger, 1988; Liu, Han, & Li, 2019; Segall, 2002, 2004; Shu, Henson, & Luecht, 2013). These models include an additional term that reifies uncertainty in the item parameters and explicitly accounts for the impact of the fluctuation when scaling the trait levels. Another relevant approach is statistical testing of descriptives of the observable variables (e.g., Choe, Zhang, & Chang, 2018; Guo, Robin, & Dorans, 2017; Lee & Lewis, 2021; Zhang, 2014; Zhang & Li, 2016; Wang & Liu, 2020). These procedures examine patterns of summary statistics of the indicator variables (e.g., proportion of correct responses) and test if the statistics indicate radical change in the manifest variables. The analysis can be performed without reference to a measurement model and thus can be applied to nonparametric settings.

We note that prior work on the drift detection was mostly designed for response data or simple item response models. When an item is modeled by multiple parameters or multiple variables, the existent methods become less applicable and lose practicability. For example, in the recalibration approach, when an item involves p parameters, one needs to concurrently examine 2^p drift statistics to probe all probable drift scenarios. Evaluating multiple statistics not only induces operational burden but it can also lead to incompatible results (e.g., flagging of an item by two statistics that examine the same parameter in opposite directions). The second approach, the direct modeling of parameter drift, suffers from the inherent overfitting problem. In real settings, parameter drift occurs only occasionally and many items tend to remain stable across the groups and over time. Using drift models in these settings overly complicates the model estimation as well as interpretation of the outcomes. If items involve multiple parameters, the complications in the model estimation will further intensify and obscure the drift evaluation. Lastly, the nonparametric procedures are not originally developed for drift detection and will likely come short of necessary power. In real assessment, item parameter drift does not typically

entail drastic change in the manifest variables, especially if the drift has confounding effects (e.g., an item decreases in difficulty but increases in guessing). When an item involves multiple parameters, it can also drift partially in a subset of parameter domains with no visible change in the indicators. The use of nonparametric procedures in these situations can face deficient power, leaving nuanced changes unnoticed.

The purpose of this study is to present a multivariate monitoring framework that overcomes the limitations of the current development. We explore statistical procedures that can jointly surveil multiple item parameters and evaluate drift directly under a presumed measurement model. We in particular take note of a recent testing trend that employs complex measurement models (e.g., three-parameter, polytomous, multidimensional, cognitively diagnostic, multi-level models) and auxiliary variables (e.g., response time, answer change, frequency in behavior). These applications typically involve multiple item parameters and require simultaneous monitoring of multiple parameter domains. As described above, using existing methods in these settings will likely induce operational drudgery and statistical problems. In this study, we explore a multivariate parametric monitoring scheme that steps over the limitations of the current methods. Our strategy is to perform direct monitoring of individual item parameters through online calibration and parametric likelihood-ratio tests. We recurrently recalibrate items and conduct generalized likelihood ratio tests to determine instance of parameter drift. All procedures are executed during test administration so that drift can be informed in real time.

Under the general scheme of multivariate parametric monitoring, we propose a number of monitoring statistics that accommodate various settings and objectives. The statistics are developed in different forms according to the sampling design, weighting, likelihood function, and summarization of drift information. Broadly, the procedures are designed to perform two sorts of monitoring: continuous and intermittent monitoring. Continuous monitoring evaluates items continuously after every use or at very frequent

time intervals. Intermittent monitoring evaluates items every once in a while when sufficient data are collected that enable new recalibration. Within each scenario, we propose a variety of decision statistics that take into account unique monitoring settings. We highlight that the monitoring procedures presented in this study demonstrate clear distinction from the existent methods. They (i) allow monitoring of multiple item parameters using singular indicating statistics, (ii) adopt powerful parametric techniques such as online calibration and likelihood-ratio tests, and (iii) perform real-time monitoring. When evaluated empirically in the simulated and real data, they showed greater sensitivity and speediness than the existing methods. The empirical evaluation suggested that the proposed monitoring schemes can be used as a powerful control tool for maintaining the quality of items.

In the following, we detail specific monitoring strategies and their empirical performance in the simulated and real data. In Section 2, we present hypotheses and online calibration technique primed for the monitoring procedures. We then discuss strategies for continuous and intermittent monitoring in Sections 3 and 4. For the former, we propose sequential generalized likelihood ratio tests (SGLRTs); for the latter, we propose cumulative sum (CUSUM) control charts. Each technique uses unique indicating statistics and sampling methods to accommodate different monitoring settings and objectives. Section 5 presents a simulation study that validated performance of the monitoring procedures under known settings. Section 6 gives an illustrative example of applying the monitoring procedures in the real assessment data. Section 7 concludes with a summary of the findings and presents some practical suggestions for operational use.

2. Online Item Monitoring

Our approach to item monitoring is to conduct parametric evaluation under the presumed measurement model. We re-estimate item parameters each time and compare new estimates with the known values. Let ξ_0 denote the parameter vector known for the studied item. At evaluation time t , the monitoring procedure recalibrates the item and

1 obtains a new parameter vector, ξ_t . The procedure then conducts a statistical test to
 2 compare the two parameter vectors:

$$3 \quad \mathcal{H}_0 : \xi_t = \xi_0 \quad vs. \quad \mathcal{H}_a : \xi_t \neq \xi_0. \quad (1)$$

4 Note that ξ_0 in (1) does not necessarily equal the true parameters. The monitoring
 5 procedure is aimed to detect digression from the known parameter values rather than from
 6 the true parameters. If ξ_0 was not accurate, the program will signal the disagreement,
 7 suggesting the need for parameter update or item re-appraisal.

8 In (1), the hypotheses are evaluated based on the estimated parameter values. The
 9 initial parameter vector, ξ_0 , can be obtained from existent data, for example, using a
 10 reference sample, or when developing an item pool (e.g., field-testing). For the new
 11 parameter vector, ξ_t , we perform online calibration to obtain brand-new parameter
 12 estimates at each evaluation time. Let $\mathcal{S}_t$ denote the calibration sample obtained at
 13 evaluation time t . The specific sampling scheme for $\mathcal{S}_t$ is determined by the monitoring
 14 procedure. Given $\mathcal{S}_t$, the new item parameter vector is obtained as the mode of the item's
 15 posterior probability distribution:

16 $\hat{\xi}_t = \arg \max_{\xi} \log \mathcal{L}(\xi | \mathbf{X}) = \arg \max_{\xi} \sum_{i \in \mathcal{S}_t} \log \iint f(\mathbf{x}_i | \boldsymbol{\eta}, \xi) f(\boldsymbol{\eta} | \boldsymbol{\Theta}_{\eta}) d\boldsymbol{\eta} + \log f(\xi | \boldsymbol{\Theta}_{\xi}),$
 17 where $\mathbf{X} = (\mathbf{x}_i : i \in \mathcal{S}_t)$ contains the calibration data, $\boldsymbol{\eta}$ is a latent variable vector, and $\boldsymbol{\Theta}_{\eta}$
 18 and $\boldsymbol{\Theta}_{\xi}$ each give the hyperparameters of the person and item parameters. Note that in
 19 real settings the calibration data may exhibit extra variation owing to test design (e.g.,
 20 adaptive item assignment) or samples (e.g., seasonality, demographics). We address this
 21 issue by using the posterior probability density, $f(\boldsymbol{\eta} | \mathbf{x}_i, \boldsymbol{\Theta}_{\eta})$, in lieu of the prior density,
 22 $f(\boldsymbol{\eta} | \boldsymbol{\Theta}_{\eta})$. This adjustment better reflects the characteristics of the calibration samples and
 23 also places the item parameter estimates readily on the scale of the item pool without
 24 requiring additional linking (Ban, Hanson, Wang, Yi, & Harris, 2001).

25 As we obtain a pair of parameter vectors, $(\hat{\xi}_0, \hat{\xi}_t)$, we conduct a statistical test to
 26 determine existence of parameter drift. A natural test that can be used in this setting is a
 27 generalized likelihood ratio test. Since the new parameter values are obtained as the

maximizer of the likelihood of $\mathcal{H}_a$, one can examine the generalized log-likelihood ratio to compare likelihoods of the known and current parameter values. Specifically, the test statistic that evaluates (1) is constructed as

$$\log \frac{\max_{\boldsymbol{\xi}} \mathcal{L}(\boldsymbol{\xi}; \mathcal{H}_a)}{\mathcal{L}(\boldsymbol{\xi}; \mathcal{H}_0)},$$

where $\mathcal{L}(\boldsymbol{\xi}; \cdot)$ gives the likelihood of $\boldsymbol{\xi}$ under the given hypothesis. A large ratio suggests that the two parameter vectors implied distinctly different likelihoods for the same indicator data in $\mathcal{S}_t$. A significantly large statistic beyond a critical value casts evidence for item parameter drift.

The use of a likelihood-ratio statistic is especially advantageous in the present setting because it can succinctly summarize the heterogeneity present across the multiple parameter domains. As will be shown below, it also allows sequential testing under various sampling designs. The following two sections discuss specific strategies for formulating the likelihood-ratio test statistics under each continuous and intermittent monitoring setting. All statistics are formulated as generalized log-likelihood ratios. They take different forms depending on the sampling design and likelihood function.

3. Continuous Monitoring

The first monitoring scenario we consider is continuous monitoring where an item is evaluated after every use. We assume that evaluation samples of the items increase by one unit each time and the monitoring procedure can utilize all data gleaned up to the evaluation point. Our strategy for this scenario is to obtain the first evaluation data and perform sequential continuous monitoring using the augmenting data. That is, the process requires preparatory downtime to initiate monitoring. Once the item is administered to a minimal admissible sample, it continuously evaluates the item after every collection of a new observation. The monitoring process continues until the item retires (e.g., maximum use) or until the item is flagged as drift.

Since the evaluation data augment over time, some sampling techniques can be used to reflect recency of the observations. In this study, we explore three sampling strategies

for conducting the SGLRTs. The first strategy is to make use of all data available at the time of evaluation. The second strategy similarly utilizes all available data but applies temporal weights to enhance sensitivity to the item's current state. The third strategy discards old observations and uses only limited sets of most recent observations in the interest of time sensitivity. Below expounds details of each method.

3.1. SGLRT

The first sampling strategy uses all observations accrued from the outset of monitoring. Let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t)$ represent the set of observations available at time t . Each element of $\mathbf{x}$ consists of multivariate cross-sectional data. The likelihood ratio statistic that evaluates parameter drift is calculated as

$$L_t = \frac{f(\mathbf{x}; \boldsymbol{\xi}_t)}{f(\mathbf{x}; \boldsymbol{\xi}_0)},$$

where $\boldsymbol{\xi}_0$ is the null item parameter vector, $\boldsymbol{\xi}_t$ is the new item parameter vector estimated from $\mathbf{x}$, and $f(\mathbf{x}; \boldsymbol{\xi})$ gives joint likelihood of $\mathbf{x}$ under $\boldsymbol{\xi}$. The functional form of f is determined by the measurement model. If the model contains latent variables, f is evaluated as marginal likelihood.

Suppose the item experienced parameter drift at time v ($1 \leq v \leq t$) with the parameter shift to $\boldsymbol{\xi}_1$. The measurement likelihood can be then re-evaluated as

$$f(\mathbf{x}; \boldsymbol{\xi}_t, v) = \left(\prod_{j=1}^{v-1} f(\mathbf{x}_j; \boldsymbol{\xi}_0) \right) \left(\prod_{j=v}^t f(\mathbf{x}_j; \boldsymbol{\xi}_t) \right).$$

The change point is one of the parameters that need be estimated by the monitoring procedure. We replace the unknown parameters (i.e., $\boldsymbol{\xi}_1$ and v) by the maximum likelihood (ML) estimates and obtain a generalized log-likelihood ratio:

$$G : G_t = \max_{1 \leq s \leq t} l_t(s) = \max_{1 \leq s \leq t} \sum_{j=s}^t \log \frac{f(\mathbf{x}_j; \hat{\boldsymbol{\xi}}_t)}{f(\mathbf{x}_j; \hat{\boldsymbol{\xi}}_0)}, \quad (2)$$

where $\hat{\boldsymbol{\xi}}_0$ and $\hat{\boldsymbol{\xi}}_t$ each represent item parameter vectors that maximize the likelihood of the null and alternative hypotheses. The statistic (2) is defined such that a large value indicates large deviation from the known values. If G_t exceeds a decision limit, h , it is concluded that $\hat{\boldsymbol{\xi}}_t$ better represents the current evaluation sample. In usual settings, h is

determined based on the asymptotic behavior of the test statistic (e.g., Lai, 1991). In this study, we obtain h via simulation to address the error inherent in the sampling and item calibration. The threshold, h , is obtained as the value that satisfies $P_{\mathcal{H}_0}(G_t \geq h) = \alpha$, where α is the significance level. If $\mathcal{H}_a$ is decided, the drift point is estimated as

$$\hat{v} = \inf \left\{ s : \max_{1 \leq s \leq t} \sum_{j=s}^t \log \frac{f(\mathbf{x}_j; \hat{\boldsymbol{\xi}}_t)}{f(\mathbf{x}_j; \hat{\boldsymbol{\xi}}_0)} > h \right\}.$$

3.2. Exponentially-Weighted SGLRT

In (2), the likelihood ratios from the different observations are treated equally. The indicating statistic can be made more sensitive to the item's current state if the ratios are weighed differentially according to the recency of data. The second sampling method applies the weighting technique to better reflect the item's latest state. Specifically, we apply the exponential weighting scheme that has been widely used and substantiated in the statistical process control (Basseville & Nikiforov, 1993; Lowry, Woodall, Champ, & Rigdon, 1992). It assigns exponentially increasing weights to newer log-likelihood ratios and exponentially decreasing weights to old ones so that the impact of old null observations can be cast out at an exponential rate. In the present setting, the exponentially-weighted monitoring statistic is constructed as

$$G^{ew} : G_t = \sum_{j=0}^t \tau_j \log \frac{f(\mathbf{x}_{t-j}; \hat{\boldsymbol{\xi}}_t)}{f(\mathbf{x}_{t-j}; \hat{\boldsymbol{\xi}}_0)}, \quad (3)$$

with τ_j weighing the j th log-likelihood ratio. The functional form of τ_j is determined by the smoothing parameter that indicates the rate of information decaying. In this study, we apply $\tau_j = \omega(1 - \omega)^j$ where $\omega \in (0, 1]$. The specific value of ω is determined in line with the monitoring settings. For example, if the monitoring is aimed to detect small shift, small ω s may be used that heavily smooth the indicating statistics. If the goal is to detect large shift quickly, large ω s are appropriate that degenerate old information at faster rate. As with (2), the significance of (3) is determined by comparing against a pre-defined decision limit— G^{ew} larger than the threshold gives evidence for parameter drift and terminates the monitoring process.

3.3. Moving-Sample Based SGLRT

Another way to conduct SGLRTs is to use moving samples that evolve over time (Choe et al., 2018; Zhang, 2014; Zhang & Li, 2016). Unlike the other two procedures, which utilize all accumulated data, the moving-sample based SGLRT applies limited sets of most recent observations. Let $\mathcal{M}_t = \{i : i = t - m + 1, \dots, t\}$ denote the moving sample obtained at time t (i.e., the group of m examinees who received the item most recently). The monitoring algorithm calibrates an item using $\mathcal{M}_t$ and compares the new parameter estimates with the values obtained from the reference sample. Let $\hat{\xi}_t$ denote the item parameter vector calibrated from $\mathcal{M}_t$, and $\hat{\xi}_0$ the parameter vector obtained from the reference sample, $\mathcal{R} = \{i : i = 1 \dots r\}$, where r denotes the reference sample size.¹ The generalized log-likelihood ratio statistic based on $\mathcal{M}_t$ is obtained as

$$G^{ms} : G_t = \sum_{i \in \mathcal{M}_t} \log \frac{f(\mathbf{x}_i; \hat{\xi}_t)}{f(\mathbf{x}_i; \hat{\xi}_0)}. \quad (4)$$

As with the other SGLRTs, (4) signals parameter drift when G^{ms} exceeds the decision limit. A drift point is estimated as the infimum of ts that satisfy $G_t > h$.

4. Intermittent Monitoring

The second monitoring scenario we consider is intermittent monitoring. This situation occurs in practice when tests are administered at time intervals or when data are collected in batches. Suppose an item has been administered to a sufficiently large sample since the last evaluation and the item can be calibrated using the newly collected data. Unlike continuous monitoring, intermittent monitoring is performed on nonoverlapping samples and an item can be calibrated independently each time. Let $\hat{\xi}_t$ represent the new parameter vector estimated from the latest sample and $\hat{\xi}_0$ denote the known parameter vector. The instance of parameter drift can be then evaluated in two ways. One is to compare measurement likelihoods of $\hat{\xi}_0$ and $\hat{\xi}_t$ under the presumed model. This strategy is akin to conducting a generalized likelihood ratio test based on the latest moving sample.

¹Previous studies defined the reference sample as $\mathcal{R} = \{i : i = 1, \dots, t - m\}$ such that it increases as the monitoring progresses. This study uses a fixed reference sample to alleviate the probable impact of false negatives in the expanding reference sample.

The other way is to compare parameter vectors directly as in the two-sample t -test. This strategy is especially appropriate in the current setting because evaluation samples are obtained as exclusive independent samples. Since the SGLRT approach was discussed above, this section focuses on the latter strategy.

Observe that in the current setting, only one pair of item parameter vectors is compared each time. For determining the occurrence of drift, it is essential to draw upon the asymptotic properties of the estimator. Let us apply a standardized distance measure that evaluates heterogeneity between the parameter vectors:

$$\boldsymbol{\gamma}_t = \boldsymbol{\Sigma}_t^{-1/2}(\hat{\boldsymbol{\xi}}_t - \hat{\boldsymbol{\xi}}_0)$$

with $\boldsymbol{\Sigma}_t = \text{Var}(\hat{\boldsymbol{\xi}}_t - \hat{\boldsymbol{\xi}}_0)$. The variance of each parameter estimate is obtained as an asymptotic estimate, i.e., the inverse of the observed Fisher information matrix. Since both the parameter vectors are obtained from the ML estimator, $\boldsymbol{\gamma}_t$ is expected to follow an approximate multivariate normal distribution:

$$\boldsymbol{\gamma}_t \overset{app}{\sim} \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (5)$$

with $\boldsymbol{\mu} = \mathbf{0}_p$ and $\boldsymbol{\Sigma} = \mathbf{I}_p$. Significant deviation of $\boldsymbol{\gamma}_t$ from $\mathbf{0}_p$ casts evidence for parameter shift. The approximate normality in (5) can be evaluated by conducting a generalized likelihood ratio test. In the null case of no significant parameter drift, a test statistic,

$$\mathcal{T}^2 = \boldsymbol{\gamma}_t^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}_t, \quad (6)$$

follows a χ^2 -distribution with p degrees of freedom when $\boldsymbol{\Sigma} = \mathbf{I}_p$ (Wilks, 1938). The test is called a generalized likelihood ratio test because $\mathcal{T}^2$ is obtained as a general statistic that maximizes the likelihood of $\mathcal{H}_a$ (i.e., by replacing the unknown $\boldsymbol{\mu}$ under $\mathcal{H}_a$ by the ML estimate). In multivariate hypothesis testing, the test is known as Hotelling's $\mathcal{T}^2$ test (Hotelling, 1931). If the test is conducted sequentially to monitor quality of a statistical process, it is called Shewhart control chart.

The Shewhart chart is of interest in the present setting because item monitoring can be seen as tracking a statistical process of an item across time. Our simulation study however suggests that application of the Shewhart chart in the current setting can result in

inadequate monitoring outcomes because of the sampling and calibration error.² One effective alternative that alleviates the oversensitivity is the CUSUM control chart (Page, 1954). Instead of charting raw fluctuations, the CUSUM chart tracks only the practically meaningful changes and signals when the cumulated evidence suggests systematic shift in the underlying process. In this study, we apply this CUSUM technique to monitor item parameter drift.

The CUSUM control chart has been used in the measurement settings to detect item compromise (e.g., DeMars, 2004; Lee & Lewis, 2021; Veerkamp & Glas, 2000; Yang, Ferdous, & Chin, 2007) and aberrant test-takers (e.g., Armstrong & Shi, 2009; Tendeiro, Meijer, Schakel, & Maij-de Meij, 2013; van Krimpen-Stoop & Meijer, 2001). Among others, Veerkamp and Glas (2000) has most relevance to the current study in that the chart was used to monitor change in the item parameters. This pioneering work however examined only a few item response parameters and assumed a particular direction when formulating the chart. When items involve multiple parameters, the existent procedure becomes operationally burdensome and can also lead to indecisive conclusions. In this study, we present multivariate alternatives that perform exploratory monitoring of multiple item parameters. The charts are developed using likelihood-ratio statistics and warrant clear decisions on the drift status. Additionally, they do not require any prior specification on the drift appearance and can be used for exploratory monitoring.

Below we give an example of a univariate chart to illustrate the key idea of a CUSUM control chart. We then present three multivariate charts that serve different functions. The first two charts monitor the location of an item (i.e., parameter values) each using a scalar and vector charting statistic. The third chart monitors both the location and dispersion (i.e., parameter values and error variance) based on a vector charting statistic. Within each

²Under the simulation design in Section 5, sequential testing based on $\mathcal{T}^2$ showed average false positive rate of 16.88%. The procedure also tended to flag drift items prematurely before the actual parameter shift, exhibiting early detection rate of 9.70%. The chart seemed overly sensitive to small fluctuations that occur from the sampling and calibration error. Note that, unlike standard Shewhart control charts, which examine manifest variables, the Shewhart chart based on $\mathcal{T}^2$ examines estimable parameters and can be influenced by the sampling and estimation error.

procedure, we explore a weighting scheme to improve time sensitivity. We note that the control charts developed in this study are intended for estimable item parameters. The control charts are originally designed for observable variables, and as alluded to above, they are not generally suitable for identifying subtle change in the item parameters. Exploiting the online calibration technique, the current study directly estimates item parameters and monitors the multiple parameter domains via multivariate CUSUM control charts.

4.1. Univariate Control Chart

Let us consider a control chart that monitors decrease in the item difficulty. Let $\hat{b}_0$ denote the difficulty level known for a studied item. At time t , a new estimate is obtained as $\hat{b}_t$ from recalibration. The monitoring procedure then evaluates difference between the two values on the standardized scale, $\eta_t = \frac{\hat{b}_0 - \hat{b}_t}{\text{SE}(\hat{b}_0 - \hat{b}_t)}$, and compares with a reference value, k , to determine significance of the difference. If $\eta_t > k$, the procedure deems the deviation noteworthy and charts the difference, $\eta_t - k$. If $\eta_t < k$, it resets the chart to zero and keeps positive shift only (i.e., decrease in difficulty). The fluctuation in the difficulties is tracked across time via a control chart:

$$C_t = \max\{0, C_{t-1} + \eta_t - k\} \quad (7)$$

with $C_0 = 0$. The monitoring continues until C_t exceeds an upper control limit, h , or until it reaches a termination point. See Figure 1 for an example scenario where an item is flagged as drift at the 23th evaluation point.

4.2. Multivariate Control Chart Based on a Scalar Charting Statistic

The univariate chart (7) is of service when a single parameter is monitored in a particular direction. When multiple item parameters need be monitored concurrently without direction, a more flexible chart must be used. In this study, we present three multivariate charts that serve different functions. The first chart monitors overall deviation in the parameters using a scalar charting statistic. Let $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$ each denote the mean vector and covariance matrix assumed for $\boldsymbol{\gamma}_t$. In the null case of no parameter drift, the

Figure 1. Univariate CUSUM Control Chart That Monitors Decrease in Difficulty

The chart tracks positive shift, $\eta_t - k$ (i.e., b -decrease larger than k) and cumulates deviations over time. The process continues until C_t exceeds the decision limit, $h = 2$, (i.e., 23th evaluation point).

parameters are expected to follow a normal distribution with $\boldsymbol{\mu}_0 = \mathbf{0}_p$ and $\boldsymbol{\Sigma}_0 = \mathbf{I}_p$. To assess likelihood of the null distribution, we apply the Mahalanobis distance measure,

$$\eta_t = \sqrt{(\boldsymbol{\gamma}_t - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\gamma}_t - \boldsymbol{\mu}_0)},$$

that transforms the vector distance into a scalar. The squared value, η_t^2 , becomes the generalized log-likelihood ratio that evaluates $\mathcal{H}_0 : \boldsymbol{\mu} = \boldsymbol{\mu}_0$ under known $\boldsymbol{\Sigma}_0$. As mentioned earlier, simple significance testing on η_t^2 (i.e., Shewhart chart) is subject to Type I error inflation. To mitigate this problem, we track only the practically significant deviations using a cumulative sum chart:

$$C^s : C_t = \max\{0, C_{t-1} + \eta_t - k\}, \quad (8)$$

where k is the reference parameter that prescribes the size of η_t worth charting. The control chart (8) records the deviations that represent practically meaningful changes. As with the regular control chart, it begins with $C_0 = 0$ and signals parameter drift when $C_t > h$.

In (8) all η_t s are treated indifferently across the evaluations. If desired, η_t s can be weighed differentially to better reflect the item's current state. Applying exponential

weights, for example, a weighted control chart can be obtained as

$$C^{s-ew} : C_t = \omega\eta_t + (1 - \omega)C_{t-1}, \quad (9)$$

where ω ($\in (0, 1]$) is a smoothing constant that weighs the latest η_t . A specific value of ω is again determined according to the monitoring design. If an item is evaluated frequently, small ω s are desirable that spread across the large number of evaluation points. If an item is seldom monitored, relatively large ω s are appropriate that span across several evaluations.

4.3. Multivariate Control Chart Based on a Vector Charting Statistic

The scalar-based control chart summarizes the deviation in the multiple parameters using a scalar statistic. In other situations, it may also be advantageous to track deviance in the individual parameters separately using a vector statistic. The vector statistic provides diagnostic information that can help identify the cause of drift while allowing a categorical decision on the drift status.

The second control chart is developed in this view. We propose a multivariate CUSUM chart that monitors multiple item parameters based on a vector-valued statistic. The chart is developed applying the Croiser's (1988) formulation.³ Croiser's approach is to use a vector charting statistic when monitoring the process behavior and convert the statistic onto the Mahalanobis metric when comparing with the scalar benchmark values (e.g., k , h). In the present setting of drift evaluation, a control chart is formulated as

$$C^v : \mathbf{C}_t = \begin{cases} \mathbf{0} & \text{if } S_t \leq k \\ (\mathbf{C}_{t-1} + \boldsymbol{\gamma}_t)(1 - k/S_t) & \text{if } S_t > k, \end{cases} \quad (10)$$

where $\mathbf{C}_0 = \mathbf{0}$ and $S_t = ((\mathbf{C}_{t-1} + \boldsymbol{\gamma}_t)^\top \boldsymbol{\Sigma}_0^{-1} (\mathbf{C}_{t-1} + \boldsymbol{\gamma}_t))^{1/2}$. The coefficient S_t transforms the vector charting statistic into a scalar so that it can be compared against the scalar reference (k) and threshold (h) values. The process signals parameter drift when the

³We note that there are other ways of constructing a multivariate chart (e.g., Healy, 1987; Pignatiello & Runger, 1990; Woodall & Ncube, 1985). These procedures however make impractical assumptions (e.g., known directions or multiple univariate charts) or make little difference in the monitoring statistics in the present setting because only one observation is evaluated each time.

Mahalanobis distance of the charting statistic, $(\mathbf{C}_t^\top \boldsymbol{\Sigma}_0^{-1} \mathbf{C}_t)^{1/2}$, exceeds a decision limit, h .

If swift drift detection is of interest, differential weights can be applied to make the charting statistic more sensitive to the current item state. Let $\boldsymbol{\Omega} = \text{diag}(\omega_1, \dots, \omega_p)$ represent the weight matrix that signifies the importance of the parameter domains. Since item parameters are evaluated individually, the parameter domains can be weighed differentially on the condition that $\sum_p \omega_p = 1$. The weighted control chart is then constructed as

$$\mathbf{C}^{v-ew} : \mathbf{C}_t = \boldsymbol{\Omega} \boldsymbol{\gamma}_t + (\mathbf{I}_p - \boldsymbol{\Omega}) \mathbf{C}_{t-1}, \quad (11)$$

where $\mathbf{I}_p$ is the $p \times p$ identity matrix. The chart begins with an initial value, $\mathbf{C}_0 = \boldsymbol{\mu}_0 = \mathbf{0}_p$, and terminates when $\mathbf{C}_t^\top \boldsymbol{\Sigma}_{C_t}^{-1} \mathbf{C}_t > h$, where $\boldsymbol{\Sigma}_{C_t}$ is covariance matrix of $\mathbf{C}_t$. In the present setting, $\boldsymbol{\Sigma}_{C_t}$ is calculated as

$$\boldsymbol{\Sigma}_{C_t} = \text{Var}(\mathbf{C}_t) = \sum_{j=1}^t \boldsymbol{\Omega} (\mathbf{I} - \boldsymbol{\Omega})^{t-j} \text{Var}(\boldsymbol{\gamma}_j) (\mathbf{I} - \boldsymbol{\Omega})^{t-j} \boldsymbol{\Omega}.$$

If the same weights are used across the domains (i.e., $\omega_1 = \dots = \omega_p = \omega$), $\boldsymbol{\Sigma}_{C_t}$ is simplified to

$$\boldsymbol{\Sigma}_{C_t} = \frac{\omega(1 - (1 - \omega)^{2t})}{2 - \omega} \text{Var}(\boldsymbol{\gamma}_t).$$

When the number of evaluation points, t , approaches ∞ , $\boldsymbol{\Sigma}_{C_t}$ is approximately obtained as

$$\boldsymbol{\Sigma}_{C_t} = \frac{\omega}{2 - \omega} \text{Var}(\boldsymbol{\gamma}_t) \quad (\text{Lowry et al., 1992}).$$

4.4. Multivariate Control Chart for Dual Monitoring

While parameter drift is generally manifested by change in the item location, increased uncertainty in the parameter estimates can also signify drift. If this is the case, evaluating both the parameter estimates and error variance can improve the monitoring performance. The third control chart is developed in light of this and monitors both the mean and variance of $\boldsymbol{\gamma}_t$. The procedure is called a dual control chart as it evaluates two aspects of an item using a single charting statistic. As with the other charts, the dual chart uses a generalized log-likelihood ratio to evaluate heterogeneity in the parameter estimates and error variance. Replacing the unknown values in $\mathcal{H}_a$ by the ML estimates, a drift

statistic is obtained as

$$g_t = \text{tr}(\mathbf{\Sigma}_t) - \log |\mathbf{\Sigma}_t| - p + \|\boldsymbol{\gamma}_t\|,$$

where $|\cdot|$ denotes the determinant of a matrix and $\|\cdot\|$ gives an L_2 norm of a vector. A large g_t implies that significant change has occurred in either the item location or error variance (i.e., $\mathcal{H}_a : \boldsymbol{\mu} \neq \mathbf{0}$ or $\mathbf{\Sigma} \neq \mathbf{I}_p$). If g_t is larger than a threshold value, the shift is considered practically significant and is documented in the control chart:

$$C^d : C_t = \max \{0, C_{t-1} + g_t - k\} \quad (12)$$

where $C_0 = 0$. The process signals compromise when $C_t > h$.

In (12), the weighting technique can be realized in two ways: by weighing (i) g_t s, or (ii) $\boldsymbol{\gamma}_t$ s and $\mathbf{\Sigma}_t$ s separately. The first method differentially weighs the current generalized log-likelihood ratio and the previous charting statistic:

$$C_t = \omega g_t + (1 - \omega)C_{t-1}.$$

The second method reweighs $\boldsymbol{\gamma}_t$ s and $\mathbf{\Sigma}_t$ s when calculating the generalized log-likelihood ratio (e.g., Zhang, Li, & Wang, 2010). Let $\mathbf{y}_t$ and $\mathbf{z}_t$ each denote the weighted form of $\boldsymbol{\gamma}_t$ and $\mathbf{\Sigma}_t$:

$$\mathbf{y}_t = \mathbf{\Omega}\boldsymbol{\gamma}_t + (\mathbf{I}_p - \mathbf{\Omega})\mathbf{y}_{t-1}$$

$$\mathbf{z}_t = \mathbf{\Omega}\mathbf{\Sigma}_t^* + (\mathbf{I}_p - \mathbf{\Omega})\mathbf{z}_{t-1},$$

with $\mathbf{\Sigma}_t^* = (\boldsymbol{\gamma}_t - \mathbf{y}_t)^\top (\boldsymbol{\gamma}_t - \mathbf{y}_t)$. A charting statistic is then constructed as

$$C^{d-ew} : g_t = \text{tr}(\mathbf{z}_t) - \log |\mathbf{z}_t| - p + \|\mathbf{y}_t\|^2, \quad (13)$$

with $g_0 = 0$. The process continues until $G_t > h$ —i.e., when there exists evidence of significant change in $\mathbf{y}_t$ or $\mathbf{z}_t$. Observe that the latter approach allows unique weight for each parameter domain. This study adopts the second approach for generality.

5. Simulation Study

A simulation study was conducted to evaluate performance of the monitoring procedures under known settings. We in particular considered a scenario an item is characterized by two indicators: the response and response time. These variables have been

used as primary indicators of test-taking process (Klein Entink, Kuhn, Hornke, & Fox, 2009) and can be immediately obtained from computer-delivered tests. They can also mirror operational functioning of items and indicate item parameter drift. Below we present simulation settings, evaluation criteria, and analysis results.

5.1. Design

Model. We applied the three-parameter logistic (3PL) response model (Birnbaum, 1968) and the log-normal response-time model (van der Linden, 2006) to model each indicator variable. The 3PL model defines the probability of a correct response as

$$P_i = \Pr(U_i = 1; \theta_i, a, b, c) = c + (1 - c) \text{logit}^{-1}(Da(\theta_i - b)), \quad (14)$$

where $U_i = u_i \in \{0, 1\}$ is a response from examinee i ; $\theta_i (\in \mathbb{R})$ is the examinee's ability level; and $a (\in \mathbb{R}^+)$, $b (\in \mathbb{R})$, and $c (\in [0, 1])$ each denote item's discrimination, difficulty, and lower-asymptote parameters. The scaling coefficient D approximates the logistic function to normal ogive. The log-normal response-time model applies the log-normal distribution to characterize a response-time variate:

$$f(T_i = t_i; \tau_i, \alpha, \beta) = \frac{\alpha}{t_i \sqrt{2\pi}} \exp \left(-\frac{\alpha^2}{2} (\log t_i - (\beta - \tau_i)^2) \right), \quad (15)$$

where t_i is an observed T_i ; $\tau_i (\in \mathbb{R})$ is examinee i 's latent speed level; and $\alpha (\in \mathbb{R}^+)$ and $\beta (\in \mathbb{R})$ are item's time-discrimination and time-intensity parameters, respectively.

The parameters of the measurement models are linked at the second level to jointly monitor the item parameters under the unified framework. We applied the van der Linden's (2007) hierarchical framework to link the parameters. The person parameters were modeled by a bivariate normal distribution, $(\theta_i, \tau_i)^\top \sim \mathcal{N}_2(\boldsymbol{\mu}_P, \boldsymbol{\Sigma}_P)$, where $\boldsymbol{\mu}_P$ and $\boldsymbol{\Sigma}_P$ are the mean vector and variance-covariance matrix of the trait variables. The item parameters were analogously linked by a multivariate normal distribution but the normality was imposed on the transformed parameters so that the parameters appear in the proper domains (e.g., $a > 0$, $\alpha > 0$, $0 \leq c < 1$). Let $\boldsymbol{\xi}$ denote the vector of transformed item parameters, $(\log a, b, \log \alpha, \beta)^\top$. The item parameters were then assumed to follow

$$\boldsymbol{\xi} = (\log a, b, \log \alpha, \beta)^\top \sim \mathcal{N}_5(\boldsymbol{\mu}_I, \boldsymbol{\Sigma}_I)$$

with $\boldsymbol{\mu}_I$ and $\boldsymbol{\Sigma}_I$ containing the mean and variance hyperparameters of the item parameters.

Item Pool. The study assumed continuous testing for evaluating the performance of the monitoring procedures. The tests were administered continuously using an item pool that consists of 100 items. The parameters of the items were generated from a quadivariate normal distribution, $(\log a, b, \log \alpha, \beta) \sim N_4(\boldsymbol{\mu}_I, \boldsymbol{\Sigma}_I)$ with $\boldsymbol{\mu}_I = (-.0431, 0, -.0431, 0)$ and

$$\boldsymbol{\Sigma}_I = \begin{pmatrix} .0862 & & & \\ .0881 & 1 & & \\ .0259 & .0881 & .0862 & \\ .0881 & .3 & .0881 & 1 \end{pmatrix}.$$

These hyperparameters correspond to the means of (1, 0, 1, 0) and standard deviations of (.3, 1, .3, 1) on the original metric and characterize typical values in the well-designed norm-referenced educational assessments. The correlation among the item parameters was uniformly set at .3 on the original scale to mimic a moderately correlated multivariate process. The c -parameter values were fixed at .2 throughout.

Initial calibration. Prior to item monitoring, we calibrated items to obtain initial item parameter values. The pre-calibration was implemented under spiraled-block linking design. Given a pool of 100 items, we created four overlapping booklets—first consisting of items from 1 to 50, second from 26 to 75, third from 51 to 100, and fourth from 76 to 100 and from 1 to 25—and administered each booklet to random samples of n_0 ($= 250, 500, 1000$) simulees. The simulees' trait levels were sampled from a bivariate normal distribution with $\boldsymbol{\mu}_P = (0, 0)^\top$ and $\boldsymbol{\Sigma}_P = \begin{pmatrix} 1 & .3 \\ .3 & 1 \end{pmatrix}$. As we obtain the item and person parameters, we generated indicator data following (14) and (15) and applied this data to estimate item parameters. The item parameters were estimated by the marginal ML estimator with Bayesian prior (Kang, Zheng, & Chang, 2020).

Continuous testing. Continuous testing was simulated by sequentially administering parallel tests to a large population of fictitious test-takers. The examinees' trait levels were drawn from the same bivariate normal distribution assumed in the

pre-calibration. For each examinee, we administered a unique set of 30 items that were randomly sampled from the pre-calibrated item pool. Items were used $N = 2000$ times across the continuous testing.

Item parameter drift. The item parameter drift was simulated assuming item leakage. We randomly sampled five items from the item pool (i.e., 5%) and introduced drift when the items were assigned to the 1001th examinee (i.e., $v = 1001$). The parameters of the drift items (i.e., (a, b, α, β)) were shifted toward the left to affect a scenario where items become easier and facile after the breach (e.g., Marianti, Fox, Avetisyan, Veldkamp, & Tijmstra, 2014; Sinharay & Johnson, 2020; van der Linden & Guo, 2008; Zopluoglu, 2019). The size of drift was varied as $\Delta = .25$ and $.50$ to model small and large drift, respectively. Note that in real settings item leakage can occur gradually or impact only small examinee groups. The monitoring performance will naturally degenerate in these settings.

Item monitoring. The study examined both the classes of monitoring procedures: the SGLRTs and the CUSUM control charts. The SGLRT-based procedures (i.e., G , G^{ew} , G^{ms}) performed continuous monitoring at frequent time intervals. The CUSUM control charts (i.e., C^s , C^v , C^d , C^{s-ew} , C^{v-ew} , C^{d-ew}) performed periodic monitoring every once in a while. Within each scenario, we additionally implemented the sequential F -test (Choe et al., 2018) and the response-based bivariate control chart (Veerkamp & Glas, 2000; C^b hereafter) to evaluate relative performance of the proposed monitoring procedures. For ease of illustration, the procedures are called by their symbolic notations in the following.

The online item monitoring was achieved as follows. The SGLRT-based procedures were initiated when items were administered to minimal samples of $n_1 = 300$, 500, or 800 examinees and performed continuous monitoring every after 100 new observations. Although SGLRTs can conduct assessment after every observation in principle, doing so induces heavy computational load while bringing little practical significance. In this study, the evaluation frequency was adjusted to achieve adequate computation but kept properly

high to allow continuous item monitoring. Under the adjusted setting, items could be evaluated up to 18, 16, and 13 times under each n_1 condition. The moving-sample based SGLRT was similarly designed to perform 18, 16, and 13 evaluations. For example, when the monitoring began at $n_1 = 500$, an item was assessed sequentially based on 16 moving samples that shift with the window size of 100 (i.e., 1-500, 101-600, ..., 1501-2000). In the case of the F -test, we used the first 250 observations as a reference sample and subsequent 251-500th, 351-600th, ..., and 1751-2000th observations as moving samples to perform 16 evaluations.

The control charts used exclusive moving samples of $n = 250, 500$, and 1000 observations to perform periodic monitoring. For example, when $n = 500$, the charts began monitoring as items were administered to 500 examinees and repeated the evaluation every time the items were assigned to a new sample of 500 examinees. The current setting resulted in maximums of 8, 4, and 2 evaluations under each sample condition.

Reference parameters. The CUSUM charts require pre-defined reference parameters (k) to track practically meaningful shifts. This study obtained the reference values from the empirical sampling distributions of the null monitoring statistics. We obtained the null testing outcomes using the uncompromised item pool and examined the distributions of the monitoring statistics to identify critical points. Specifically, we identified three critical values that represent small, moderate, and large deviation for each chart and used these values as reference parameters. In the control charts that monitor item location, $k = .5, 1.0$, and 2.0 were used as reference values. In the dual chart that examines both the item location and error variance, we applied $k = 5, 10$, and 15 .

Smoothing parameters. The study also examined the impact of applying the weighing technique in each procedure. The smoothing parameters (ω) were decided in line with the monitoring settings. In SGLRTs, we applied small values ($\omega = .001, .003, \dots, .009$) to accommodate large numbers of readings (i.e., minimum of 250). In control charts, relatively large smoothing constants were used ($\omega = .1, .3, \dots, .9$)

because of the low evaluation frequency. For those charts that permit differential weights across the parameter domains (e.g., C^{v-ew} , C^{d-ew}), we applied the same weights across the domains to make them comparable to other procedures.

Decision limits. All monitoring procedures apply thresholds (h) for determining the significance of parameter shift. The current study obtained the decision limits through Monte Carlo simulation of null tests. We defined familywise Type I error, α (set at .05 throughout), as the probability of flagging an item among the evaluated items.⁴ We then constructed an empirical sampling distribution by drawing the maximum charting statistics of the items and determined the threshold values as the $(1 - \alpha)$ th quantiles of the sampling distributions.

Replication. All simulation conditions were replicated 100 times, each with a unique set of items and a sample.

Remark. We note that the simulation study was aimed to evaluate empirical performance of the monitoring procedures in an example testing scenario and some simplifications were made to keep the scope of work manageable (e.g., a certain drift scenario, fixing c at a constant). Assuming more complex settings (e.g., adaptive testing, complex models) will naturally necessitate larger samples and computation time. Although the current simulation was carried out in the controlled settings, we anticipate that the monitoring procedures would yield similar relative performance in other settings. Our additional simulations under other settings suggested that, although the absolute performance of the procedures can differ by the monitoring settings, the relative performance generally remains consistent.

5.2. Evaluation

The performance of the monitoring procedures was evaluated by four criteria: (i) Type I error rate, (ii) detection power, (iii) average run length (ARL), and (iv) premature signal. The Type I error rate evaluates the null behavior of the monitoring procedures. It

⁴In sequential testing, Type I error can be defined in three ways—across the event times, the items, and across both the events and items.

was defined as the proportion of non-drift items that were erroneously flagged as drift. The study examined the error rate among the null items that were administered with the drift items (i.e., tests created with the contaminated item pools). Since drift items adversely affect online calibration and drift evaluation, we anticipate that the monitoring procedures would display inflated error rate in flagging the null items. With this point in view, the next section pays close attention to the magnitude of Type I error inflation.

The other criteria evaluate performance on the drift items. The power and ARL evaluate effectiveness of the monitoring procedures in identifying the drift items. The early detection rate evaluates false alarm in the drift items. The power rate was defined as the correction identification of a drift item after the actual parameter shift. We applied the Cohen’s criterion (Cohen, 1992) to determine excellent ($\geq .80$) and moderate ($[\text{.60}, .80)$) power rate. The timeliness was evaluated by lag between the drift point and detection time. In the current setting, SGLRTs can have a minimum of zero and maximum of nine lags. Control charts can have a minimum of zero and maximum of one or three lag(s) depending on the number of evaluations. The zero lag means that the monitoring procedure identified the drift at the earliest possible time. The positive lags mean that there was delay in identifying the drift. Note that in sequential testing, the timeliness of detection is commonly evaluated by average run length—i.e., the average number of samples observed before a false alarm (called in-control ARL) or before a correct detection (called out-of-control ARL). The current study treated all false alarms in the drift items as premature identification and evaluated such instances using a separate criterion. That is, whenever a drift item was flagged before the actual development of parameter drift, it was counted as early detection and excluded from the power.

5.3. Results

Below presents simulation results for each evaluation criterion. Since monitoring was performed in two scenarios using unique procedures, the results are presented separately for each scenario. Where appropriate, we also present significance test results from the

variance analysis to draw distinction among the evaluated procedures.

5.3.1. Continuous Monitoring. The following evaluates four procedures, F , G , G^{ew} , G^{ms} , that performed continuous monitoring.

Type I error. The SGLRT-based procedures overall showed adequate performance in regulating the Type I error. The error rates remained close to the nominal level and showed little variation across the simulated conditions. We could not identify any patent pattern in the error rates; we therefore report average statistics without tabulation. The three SGLRT-based procedures yielded average rejection rates of .048 (G), .048 (G^{ew}), and .049 (G^{ms}). The nonparametric F -test showed slightly inflated error rate, producing .055 on average. We restate that the present results were obtained from the compromised tests that included drift items. Despite the negative bearing, the parametric monitoring procedures demonstrated robust performance in general.

Power. Table 1 reports power rates of the SGLRTs and the F -test. The results of G^{ew} were presented for three strata that represent small, medium, and large smoothing. The results from Table 1 suggest that the SGLRTs performed competently well in identifying the drift items. Despite the small drift size and small evaluation samples, the procedures constantly produced excellent power—.902 (G), .934 (G^{ew}), and .965 (G^{ms}) on average. The moving samples and exponential weights were shown to improve detection performance. Compared to G , G^{ms} and G^{ew} led to significantly higher power rates ($p < .001$) and delivered more stable performance against the variation in the design variables. Comparison between the SGLRTs and the F -test suggests that parametric monitoring indeed achieved higher power. Throughout the evaluations, the F -test showed constantly lower detection power ($p < .001$; .645 on average) and at times suffered from insufficient power especially when drift occurred to a small degree.

In Table 1, some anticipated patterns can be immediately verified. Increase in the drift size and evaluation sample led to higher power in all procedures (average power = (.828, .891, .937) when $n_1 = (300, 500, 1000)$, , $p < .001$; (.802, .969) when $\Delta = (.25, .50)$,

Table 1

Correct Detection Rates of SGLRTs

n_1	Δ	F	G	$G^{ew}(\omega)$			G^{ms}
				.001	.005	.009	
300	.25	.258	.686	.810	.792	.788	.914
	.50	.800	.988	.986	.972	.968	.976
500	.25	.392	.836	.906	.906	.882	.952
	.50	.928	.988	.980	.976	.970	.980
800	.25	.496	.916	.976	.978	.970	.982
	.50	.994	.996	.980	.984	.980	.988

Note. n_1 : Sample size at the first evaluation. Items could be evaluated 18, 16, and 13 times at maximum under each n_1 condition. Δ : Drift size. All four parameters (a, b, α, β) were shifted to the left. ω : Smoothing parameter. G : Regular SGLRT. G^{ew} : Exponentially-weighted SGLRT. G^{ms} : Moving-sample based SGLRT.

$p < .001$). In G^{ew} s, the power rates declined in reverse to the smoothing constants ((.940, .935, .926) when $\omega = (.001, .003, .005)$, $p = .133$). This tendency was expected because small ω s were indeed intended to detect small shift (i.e., greater sensitivity). As we examine ARLs below, we confirm that large ω expectedly led to faster detection of large parameter shift.

Lag. Table 2 presents out-of-control ARLs of the SGLRTs and F -test. The results show that the procedures adequately identified the drift items in a timely manner. Among the four procedures, G^{ms} and G^{ew} achieved the most rapid detection, each producing 1.341 and 1.500 ARLs. G and F identified the drift relatively slowly, yielding 2.870 and 2.882 ARLs. The significantly shorter ARLs in G^{ms} and G^{ew} ($p < .001$) suggest that the moving-sample and weighting techniques effectively expedited the drift detection. Between G^{ms} and G^{ew} , G^{ms} appeared more effective in improving the speed ($p < .01$). The performance between G and F appeared indifferent ($p = .418$).

It is worth noting that in Table 2, both F and G^{ms} showed longer ARLs as the moving samples enlarged. This is due to different proportions of the null and drift observations. For instance, when the moving sample size (n_m) is conditioned at 300, the impact of parameter drift begins to appear at the 9th evaluation with the sample containing 100 observations from the drift items and 200 observations from the non-drift

Table 2

Average Run Length of Drift Items Correctly Flagged by SGLRTs

n_1	Δ	F	G	$G^{ew}(\omega)$			G^{ms}
				.001	.005	.009	
300	.25	3.29	4.66	3.58	2.53	2.58	1.81
	.50	2.20	2.09	1.05	.52	.48	.35
500	.25	3.55	4.34	3.19	2.41	2.36	2.21
	.50	2.00	1.46	.78	.37	.31	.56
800	.25	4.38	3.73	2.64	1.76	1.70	2.54
	.50	1.88	.94	.47	.15	.11	.58

Note. n_1 : Sample size at the first evaluation. Items could be evaluated 18, 16, and 13 times at maximum under each n_1 condition. The maximum possible ARL equals 9 in the current simulation design. Δ : Drift size. All four parameters (a, b, α, β) were shifted to the left. ω : Smoothing parameter. G : Regular SGLRT. G^{ew} : Exponentially-weighted SGLRT. G^{ms} : Moving-sample based SGLRT.

items (recall that continuous monitoring was performed at every 100 new observations). As the monitoring progresses, the moving samples gradually cast out the null observations, and after two subsequent evaluations (i.e., at the 11th evaluation), the samples no longer include the null observations, fully embodying the effect of parameter shift. When $n_m = 800$, the null observations diminish at slower rate. The impact of parameter drift is materialized at the 4th evaluation point with the sample containing 100 observations from the drift items and 700 null observations, and only after seven additional evaluations, the samples no longer include the null observations. The prolonged existence of the null observations in the large moving samples confounds the impact of parameter drift and decelerates the identification. For this reason, the large moving samples tended to entail longer ARLs in Table 2.

The other patterns relating to the design variables were consistent with the expectations. The greater the parameter shift or the larger the evaluation sample, the faster the drift detection. In G^{ew} , large ω s expectedly led to shorter ARLs and faster detection.

Premature detection. Table 3 reports early detection rates of the SGLRT-related procedures. The monitoring procedures on the whole demonstrated good timeliness, maintaining false alarms within a tolerable range. Among the four procedures evaluated, G

Table 3
Early Detection Rates of SGLRTs (%)

n_1	Δ	F	G	$G^{ew} (\omega)$			G^{ms}
				.001	.005	.009	
300	.25	7.83	.50	2.48	4.57	4.62	2.40
	.50	.90	.60	1.00	2.40	2.80	2.20
500	.25	5.50	.45	2.50	3.60	3.93	2.80
	.50	.85	1.00	2.00	2.40	3.00	2.00
800	.25	1.58	.80	1.20	1.20	1.20	1.60
	.50	.40	.40	2.00	1.60	2.00	1.20

Note. n_1 : Sample size at the first evaluation. Δ : Drift size. All four parameters (a, b, α, β) were shifted to the left. ω : Smoothing parameter. G : Regular SGLRT. G^{ew} : Exponentially-weighted SGLRT. G^{ms} : Moving-sample based SGLRT.

showed the most timely behavior, with early detection rate .63% on average. Followed by were G^{ms} (2.03%), G^{ew} (2.47%), and F (2.84%), showing comparable performance. Patterns relating to the design variables were consistent with expectations. Increase in the drift size and the evaluation sample size resulted in fewer erroneous detections (average error rate = (2.69%, 2.50%, 1.27%) when $n_1 = (300, 500, 800)$, $p < .001$; (2.71%, 1.60%) when $\Delta = (.25, .50)$, $p < .001$). Using smaller smoothing coefficients in G^{ew} led to lower error rates though the differences were generally marginal ($p = .156$).

5.3.2. Periodic Monitoring. Below we present results of the control charts (i.e., C^b, C^s, C^v, C^d) that performed periodic monitoring. The results of the weighted charts were omitted because they showed similar performance with the unweighted charts.

Type I error. The control charts showed acceptable performance in regulating the Type I error. We report average false rejection rates without tabulation. The four charts under evaluation produced average rejection rates of .044 (C^b), .053 (C^s), .051 (C^v), and .052 (C^d). As mentioned earlier, the inflation in the error was expected because the results were obtained from the compromised tests that include drift items. Notwithstanding, the procedures seemed to maintain close adherence to the nominal level. Among the control charts, the multivariate charts appeared more prone to error. This is because the procedures were implemented with little information. Recall that C^b was performed

applying correct knowledge about the drift appearance (e.g., the type of parameters drifted, direction of shift). When C^b was implemented under incorrect priors, it showed higher false rejection rates and near-zero detection power.

Power. Table 4 reports correct detection rates of the control charts. The procedures performed adequately well in identifying the parameter drift. The bivariate chart that examined ab -drift showed moderate power of .628. The quadvariate charts that examined all four parameters demonstrated excellent power—.992 (C^s), .975 (C^v), and .993 (C^d). Among the multivariate charts, C^d and C^s achieved clear outperformance over C^v ($p < .001$). Although C^v showed somewhat lower detecting power, the overall power rates were still within the excellent range. We also note that C^v may lend better utility in applied settings. Unlike the other charts, which give only summary information, C^v provides more formative information about drift status (e.g., parameter type, degree of drift) and can be more instrumental in conducting cause analysis.

Table 4 also shows some distinct patterns in relation to the design variables. As expected, increase in the drift size and the evaluation sample led to enhanced power (.855, .903, .933) when $n_1 = (250, 500, 1000)$, $p < .001$; (.832, .962) when $\Delta = (.25, .50)$, $p < .001$). Increase in the reference parameter (k) entailed mixed results. Recall that the reference parameter was designed to indicate the amount of change worthy of note. Large

Table 4
Correct Detection Rates of CUSUM Control Charts

n_1	$\Delta \backslash k$	C^b			C^s			C^v			C^d		
		.5	1	2	.5	1	2	.5	1	2	5	10	15
250	.25	.256	.292	.358	.958	.958	.960	.834	.884	.946	.968	.966	.964
	.50	.676	.744	.800	1	1	.994	.998	.996	.988	1	.998	.992
500	.25	.304	.336	.454	1	1	.996	.960	.980	.988	1	1	1
	.50	.856	.884	.916	1	1	.998	1	1	.996	1	.998	.998
1000	.25	.492	.514	.596	1	1	.998	.990	1	.998	1	1	.998
	.50	.928	.942	.948	1	1	.998	1	.998	.994	1	.998	.998

Note. n_1 : Sample size at the first evaluation. Δ : Drift size. All four parameters (a, b, α, β) were shifted to the left. k : Reference parameter that decides practically meaningful change. C^b : Response-based bivariate control chart that monitors decrease in (a, b). C^s : Scalar-based multivariate control chart. C^v : Vector-valued multivariate control chart. C^d : Dual control chart.

k s give greater assurance on the parameter shift and are thus expected to lead to higher detecting power. The results in Table 4 generally follow this expected pattern. As k increased, the power rates increased. When the control charts achieved near perfect power, however, a reverse trend appeared. As k increased, the power rates declined slightly. The latter pattern can be explained by the way the power was calculated. As described above, when a monitoring procedure flags a drift item before the actual change, we treated this instance as premature detection and counted out from the power calculation. It appeared that when a control chart employs a large reference value, the decision limit from the empirical sampling distribution tended to take a small value and it consequently makes the chart more sensitive to the small fluctuations.⁵ The increased sensitivity induced by the small decision limit made the chart more susceptible to small fluctuations and gave false alarms even before the development of drift. We note that the untimely detection of the control charts occurred only rarely and did not appear as critical practical concern. On the whole, the multivariate charts seemed to adequately identify drift items with excellent power without serious false signals.

Lag. Table 5 reports ARLs of the control charts. The results from $n_1 = 1000$ were omitted because all ARLs equaled zero. The reported values suggest that the control charts demonstrated adequate timeliness in drift detection. All procedures identified the drift far before they reach the maximum lags (3 when $n_1 = 250$; 1 when $n_1 = 500$). Comparison between the bivariate and quadivariate charts suggests that the latter led to much quicker detection (on average 1.158 (bivariate), .551 (multivariate) ; $p < .001$). Among the multivariate charts, C^d identified the drift most rapidly (.284), followed by C^s (.614) and C^v (.754) (all pairwise comparisons were significant at $p < .001$). C^d also performed most stably, showing least variation across the design variables (SD = .427).

The impact of the design variables was generally consistent with the expectations. The large drift size and evaluation samples led to faster drift detection ((.970, .435) when

⁵Recall that the charting statistics are obtained by subtracting the reference values (k). The larger the k , the smaller the null charting statistics, and thus, the smaller the decision limit.

Table 5

Average Run Length of Drift Items Correctly Flagged by CUSUM Control Charts

n_1	$\Delta \backslash k$	C^b			C^s			C^v			C^d		
		.5	1	2	.5	1	2	.5	1	2	5	10	15
250	.25	2.08	1.97	1.52	1.81	1.71	1.39	2.10	1.93	1.42	1.21	1.07	.89
	.50	2.07	1.79	.95	.76	.65	.30	1.07	.88	.34	.06	.04	.03
500	.25	.74	.71	.61	.35	.28	.14	.59	.48	.19	.05	.04	.03
	.50	.67	.55	.24	0	0	0	.04	.02	.00	0	0	0

Note. n_1 : Sample size at the first evaluation. The maximum possible ARLs under each n_1 equals 3 ($n_1 = 250$) and 1 ($n_1 = 500$). When $n_1 = 1000$, the maximum possible ARL equals zero. Δ : Drift size. All four parameters (a, b, α, β) were shifted to the left. k : Reference parameter that decides practically meaningful change. C^b : Response-based bivariate control chart that monitors decrease in (a, b). C^s : Scalar-based multivariate control chart. C^v : Vector-valued multivariate control chart. C^d : Dual control chart.

1 $\Delta = (.25, .50)$, $p < .001$; (1.167, .238) when $n_1 = (250, 500)$, $p < .001$). The use of large
2 reference values likewise resulted in quicker identification by virtue of small decision limits
3 ($p < .001$).

4 **Premature detection.** The control charts showed adequate timeliness in signaling
5 the parameter drift. The early detection rates equaled 1.28% (C^b), .11% (C^s), .19% (C^v),
6 and .14% (C^d) on average. Among the four procedures evaluated, C^b made false alarms
7 most frequently ($p < .001$) while other charts performed alike ($p > .529$). The design
8 variables had marginal impact in general. Only the reference parameter was found to have
9 significant impact ($p < .001$) with large k s entailing larger error rates.

10 **Remark on the complete and partial drift.** We note that the present results
11 were obtained when the items drift in all parameter domains. In real settings, items can
12 drift partially in a subset of parameter domains, and the performance of the monitoring
13 procedures can degenerate in line with the drift effect size. For example, our additional
14 simulation study suggests that when items drift in (a, b) parameters at the size of $\Delta = .5$,
15 average power rates of SGLRTs decline to .319 (G), .404 (G^{ew}), and .563 (G^{ms}), and those
16 of control charts drop to .498 (C^s), .391 (C^v), and .571 (C^d). While the procedures
17 regained power as they use larger samples and center on the relevant domains, the overall
18 power rates were lower than those reported above. This experimentation suggests that

multivariate monitoring is most effective when all evaluated parameter domains are subject to drift. In real settings, one may explore probable drift using the multivariate procedures and obtain supporting information using the simpler procedures. We apply this strategy when analyzing the real data below.

6. Application to Real Data

Based on the results from the simulation study, we applied the monitoring procedures to real-assessment data to examine relevance to applied settings.

6.1. Data

The data were obtained from a credential assessment program in the U.S. The data set contained 60,718 examinees' responses and response times collected over a quarter. During testing, examinees received unique sets of items that were adapted to their ability levels. The tests were drawn from a fixed pool of 1,472 items and administered at varying lengths. Across the time the item pool was in use, the items were used 4032.34 times on average ($SD = 3287.87$) with a maximum of 11,284 and a minimum of 8.

6.2. Analysis

Initial item calibration. We obtained initial item parameter values by calibrating the items prior to monitoring. The initial calibration samples were drawn from the 1000 reference test-takers (i.e., the first-time attempters) who received the items at the earliest dates. The item parameters were estimated jointly under the two-parameter logistic response model and the log-normal response-time model through marginal maximum a posteriori. As with the simulation, the estimation applied the posterior density in lieu of the prior density in order to redress the central tendency of the calibration samples toward the item difficulties—the criterion on which the items were administered. The initial item calibration resulted in a battery of 814 items that converged within the tolerance criteria. The ensuing analysis conducted sequential monitoring of these items to examine existence of parameter drift. Although the analysis was performed a posteriori, the monitoring was implemented mimicking continuous online testing. Once the monitoring was set off, we

1 calibrated the items online and evaluated parameter drift in real time.

2 **Monitoring procedures.** The performance of the items was monitored both
 3 continuously and periodically. For continuous monitoring, we applied the three SGLRTs
 4 (i.e., G , G^{ew} , G^{ms}) and the F -test. The periodic item monitoring was performed applying
 5 the control charts, including the regular multivariate charts (C^s , C^v , C^d) and their
 6 weighted variants (C^{s-ew} , C^{v-ew} , C^{d-ew}). Within each scenario, we additionally applied
 7 the procedures that examine partial drift in the response or response-time parameters. For
 8 example, in the control-chart approach, we implemented three types of control charts: one
 9 that monitors all item parameters (a , b , α , β) and two others that monitor each response
 10 (a , b) and response-time (α , β) parameters only. Where appropriate, we also applied the
 11 univariate procedures that monitor a specific item parameter in a particular direction. The
 12 SGLRTs and F -test were similarly reduced for each measurement model or for each
 13 indicator variable (e.g., Choe et al., 2018; Zhang, 2014).

14 **Design parameters.** The design parameters of the monitoring procedures were
 15 determined in line with the monitoring settings. The smoothing parameters in G^{ew} were
 16 set at small values, $\omega = (.1, .3, .5, .7, .9) \times 10^{-3}$, to accommodate large evaluation
 17 samples. Those of the control charts were set as $(.1, .3, .5, .7, .9) \times 10^{-1}$ considering the
 18 number of evaluation points. The reference parameters in the control charts were varied
 19 between .25 to 3 with increment of .25.

20 **Sequential monitoring.** The online item monitoring was implemented as follows.
 21 All procedures began monitoring when items were assigned to minimal samples of 1000
 22 examinees and performed sequential monitoring thereafter. The SGLRT-related procedures
 23 (including the F - and z -tests) evaluated items continuously as the items were administered
 24 to 100 new additional examinees. Depending on the usage, an item could be evaluated up
 25 to 103 times in continuous monitoring. In the control charts, items were evaluated
 26 periodically each time they were assigned to a new sample of 1000 examinees (i.e., window
 27 size = 1000) and could be evaluated up to eleven times depending on the usage.

Decision limits. Significance of parameter drift was determined by comparing the charting statistics with the decision limits. We obtained an empirical sampling distribution of the drift statistics by bootstrapping 3000 resamples of the observed values and determined the decision limit as the average of quantiles.⁶ For the two-sided tests (e.g., z -test), the threshold values were obtained as the averages of each 2.5th and 97.5th quantiles of the resamples. Those for one-sided tests were obtained as the average of the 95th quantiles of the sampling distributions.

Given a large number of monitoring procedures evaluated, we determined the instance of parameter drift by considering multiple procedures collectively. Specifically, we flagged an item as drift when the item was flagged at least once under the different design parameter values (e.g., reference parameters, smoothing constants) within the same monitoring procedure (e.g., C^s). The flagged items were then subsequently examined in greater depth to determine the specific drift pattern.

6.3. Results

Table 6 reports the numbers of items flagged by the monitoring procedures. The results were categorized by the measurement models to highlight different kinds of the indicator sources. We note that the results presented in Table 6 do not generally conform to the nominal level. This is because drift was determined leniently (i.e., considered drift if flagged at least once under the different design levels). Despite the liberal flagging, the drift item sets were generally small in size and they mostly overlapped across the different procedures (e.g., among the items flagged by the joint procedures, 31 items were the same items). Based on the results from Table 6, we conducted thorough review of the individual items to investigate specific drift patterns. Below we present two example items that were found to drift in (a, b, α, β) and (a, b) .

Figure 2 presents time-series plots of an item drifted in (a, b, α, β) . The first four

⁶We also contemplated simulation for attaining the threshold values. The resulting values however did not generally accord with the statistics in the real data possibly due to disparity in sampling (e.g., content-balancing and item exposure control in real testing.)

Table 6
Number of Items Flagged in Real Data

Method	SGLRT			Method	Control Chart		
	JM	IRM	RTM		JM	IRM	RTM
G	41	39	40	C^s	58	56	50
G^{ew}	59	57	53	C^v	75	68	68
G^{ms}	41	41	41	C^d	59	53	53
				C^{s-ew}	70	65	69
				C^{v-ew}	77	65	62
				C^{d-ew}	89	85	84

Note. JM: Joint model. IRM: Item response model. RTM: Response time model. G Regular SGLRT. G^{ew} : Exponentially-weighted SGLRT. G^{ms} : Moving-sample based SGLRT. C^s : Scalar-based multivariate control chart. C^v : Vector-valued multivariate control chart. C^d : Dual control chart. C^{s-ew} : Scalar-based control chart applying exponential weights. C^{v-ew} : Vector-valued control chart applying exponential weights. C^{d-ew} : Dual control chart applying exponential weights. The total number of items monitored was 814. For the procedures that involve effect sizes or smoothing constants, the union of the flagged items was considered.

plots give parameter estimates; the last three plots report average absolute difference between the initial and re-estimated parameter values, proportion correct, and average response times observed at each evaluation point. The plots were presented for the three procedures that entailed new calibration at each evaluation. The item under study was administered to 9,928 examinees and monitored 90 times by the SGLRTs and 9 times by the control charts. In normal settings, the monitoring process terminates when it senses significant parameter drift. In this study, we enforced the monitoring to continue until the end of item assignment in order to examine the functioning of the item across the entire lifespan. The time-series plots show that the item began with moderate discriminating power and medium difficulty, but became gradually easier over time and faltered in discrimination. The response-time parameters similarly showed fluctuations. The time-discriminating parameters shifted gradually to the right, suggesting greater differentiation among the fast and slow test-takers. The time-intensity parameters decreased slightly, suggesting that examinees who received the item at the later stage tended to spend less time than those who received at the earlier stage.

Table 7 reports the time points at which the monitoring procedures signaled

parameter drift. When the procedures were implemented under different design levels, we reported the earliest alarm points. The two monitoring techniques indicated similar patterns. The SGLRTs suggested $8800 \sim 8900$ (G), $7700 \sim 7800$ (G^{ew}), and $4900 \sim 5000$ (G^{ms}) as the potential drift window. The control charts suggested $4000 \sim 5000$ or $7000 \sim 8000$ as the potential drift interval. As was the case in the simulation study, the procedures that applied the exponential weights and moving samples pointed to earlier times as the drift point.

Table 7 reveals some distinct patterns in relation to the indicator variables. The joint monitoring procedures that used both the response and response-time data tended to signal parameter drift at earlier times than those response and response-time based procedures. The procedures evaluating (a, b) -drift occasionally failed to flag the item despite the apparent change in b .

The univariate and bivariate control charts indicated similar patterns. Among the eight univariate charts evaluated, the charts that evaluated decrease in b , increase in α , and decrease in β signaled drift at the 9th, 8th, and 8th evaluation points, respectively. The exponentially-weighted univariate charts suggested the same results. The bivariate control

Table 7
Estimated Drift Points in Real Data

Method	SGLRT			Method	Control Chart		
	JM	IRM	RTM		JM	IRM	RTM
G	80	-	83	C^s	8	-	8
G^{ew}	69	85	78	C^v	8	-	8
G^{ms}	41	76	69	C^d	5	-	6
				C^{s-ew}	5	-	5
				C^{v-ew}	5	-	5
				C^{d-ew}	8	-	8

Note. JM: Joint model. IRM: Item response model. RTM: Response time model. G Regular SGLRT. G^{ew} : Exponentially-weighted SGLRT. G^{ms} : Moving-sample based SGLRT. C^s : Scalar-based multivariate control chart. C^v : Vector-valued multivariate control chart. C^d : Dual control chart. C^{s-ew} : Scalar-based control chart applying exponential weights. C^{v-ew} : Vector-valued control chart applying exponential weights. C^{d-ew} : Dual control chart applying exponential weights. For the procedures that involve effect sizes or smoothing constants, the earliest alarm point was reported. The unit of analysis was 100 in SGLRTs and 1000 in control charts.

charts suggested drift in (a, b) and (α, β) , each in the direction of $(-, -)$ and $(+, -)$. The (a, b) -drift was signaled at the 9th evaluation (by C^b), and the (α, β) drift was signaled at the 9th (C^b) and 5th evaluation points (C^{b-ew}), respectively. Note that in Figure 2(a) a -estimates tended to fluctuate rather than to decrease. It seemed that the (a, b) -chart flagged the item chiefly because of the drift in b and not really due to any substantive change in a .

Figure 3 presents a second example item that was flagged as (a, b) -drift. The item was administered to 11,155 examinees and evaluated 102 times by the SGLRTs and 11 times by the control charts. The plots suggest that the item remained fairly stable in (α, β) but fluctuated a bit in (a, b) . When evaluated by the monitoring procedures, several joint procedures and all response-based procedures flagged the item. The response-based procedures signaled drift relatively at earlier times ($C^s, C^v, C^{s-ew}, C^{d-ew}$ at 8; C^d and C^{v-em} at 9; G at 93; G^{ms} at 83; G^{ew} at 82), suggesting between the 7000th and 8000th examinees (control charts) or between the 8900th and 9200th examinees (SGLRTs) as a potential drift point. The joint procedures flagged the item at later times (C^s at time 10, C^v and C^d at 11), suggesting the window between 9000th and 10,000th examinees as a potential drift interval. Among the univariate and bivariate control charts, the univariate chart signaled b -decrease at the evaluation points 4 (C) and 5 (C^{ew}). The bivariate charts gave confounding results, suggesting that (i) both the (a, b) parameters decreased (C and C^{ew} at 8) or (ii) a increased while b decreased (C at 5; C^{ew} at 4).

The present examples suggest that the monitoring procedures must be used with care. While multivariate joint monitoring can improve power and timeliness, the mechanical use can result in loss of power and slower detection when an item is drifted partially in the subset of parameter domains. The bivariate control charts, on the other hand, though useful when an item is monitored with a prescribed direction, can lead to inconsistent conclusions when no specific direction is predicated. Our simulation study and real data analysis suggest that in practice one may want to appraise an item holistically

using the multivariate procedures and then scope out specific parameter sets or individual parameters to obtain detailed information. Such a top-down approach can gradually reveal greater information about drift specifics (e.g., type, degree, direction) while allowing decisive conclusions on the presence of drift.

7. Discussion

As many testing programs use item pools these days, a quality control tool that monitors functioning of items has been greatly needed. This study is presented in an effort to provide flexible and efficient online item monitoring tools for continuous testing. We proposed multivariate parametric monitoring procedures that simultaneously examine multiple item parameter domains. All proposed procedures performed online monitoring so that drift in the item parameters can be informed in real time during test administration. Specifically, two techniques were proposed for online monitoring: SGLRTs for continuous monitoring and control charts for intermittent monitoring. The SGLRTs evaluate items continually as items are administered to new examinees. The control charts conduct monitoring once in a while based on a new batch of data. Both the techniques drew on the likelihood-ratio statistics to determine appearance of parameter drift. The SGLRTs examined measurement likelihood of the item parameters; the control charts evaluated likelihood of normality of a distance statistic. For each technique, we contemplated various monitoring statistics to accommodate different monitoring settings and objectives (e.g., availability of data, sensitivity to recent data, how to summarize information across multiple domains).

Empirical experimentation on the simulated and real data suggested that the proposed procedures perform effectively well in identifying the parameter drift. Compared with the existent methods (e.g., Choe et al., 2018; Veerkamp & Glas, 2000; Zhang, 2014), the present procedures achieved higher power and faster identification of drift items. Despite the greater complexity and uncertainty in item evaluation, the suggested methods appeared to bring clear benefit to item monitoring.

1 While our study has shown great promise for multivariate monitoring, it also
2 suggested that the procedures must be used with prudence. In particular, when items drift
3 partially in the parameter domains, the multivariate procedures can show degraded
4 monitoring performance. In real settings, item monitoring is typically performed in an
5 exploratory fashion, and it is sensible to examine multiple procedures concurrently. For
6 example, one may apply multivariate procedures to conduct a holistic assessment and use
7 reduced procedures for detailed review. Such a top-down approach can gradually reveal
8 greater information about drift appearance while allowing clear decisions on the existence
9 of drift. We note that the procedures that are reduced to specific parameter domains still
10 maintain practical advantages over the existing methods. They use singular indicating
11 statistics without assuming a particular drift direction. They also provide detailed
12 information about drift profiles that can help inform cause analysis.

13 We conclude with some limitations of the study. We mention that our simulation
14 study focused on a few main factors to maintain the discussion adequately focused. The
15 monitoring procedures can however be affected by other variables, such as the proportion of
16 drift items, drift time points, parameter type, etc. The change in these variables naturally
17 leads to different monitoring outcomes. Our additional studies suggest that, although the
18 absolute performance may vary depending on the monitoring settings, the relative
19 performance of the monitoring procedures tends to remain consistent. We also add that
20 the present study assumed homogeneous samples throughout the empirical evaluations.
21 While we attempted to address the sampling variation during item calibration, a future
22 study may explore methods that can explicitly address the seasonality in drift evaluation.

References

- Armstrong, R. D., & Shi, M. (2009). A parametric cumulative sum statistic for person fit. *Applied Psychological Measurement*, 33, 391–410.
- Ban, J. C., Hanson, B. A., Wang, T., Yi, Q., & Harris, D. J. (2001). A comparative study of on-line pretest item-calibration/scaling methods in computerized adaptive testing. *Journal of Educational Measurement*, 38(3), 191–212.
- Basseville, M., & Nikiforov, I. V. (1993). *Detection of abrupt changes: Theory and applications*. Upper Saddle River, New Jersey: Prentice-Hall, Inc.
- Birnbaum, A. (1968). Theories of mental test scores. In F. M. Lord & M. R. Novick (Eds.), *Some latent trait models and their use in inferring an examinee's ability* (pp. 397–479). MA: Addison-Wesley, Reading.
- Bock, R., Muraki, E., & Pfeifferberger, W. (1988). Item pool maintenance in the presence of item parameter drift. *Journal of Educational Measurement*, 25, 275–285.
- Choe, E. M., Zhang, J., & Chang, H.-H. (2018). Sequential detection of compromised items using response times in computerized adaptive testing. *Psychometrika*, 83, 650–673.
- Clark, A. (2013). *Review of parameter drift methodology and implications for operational testing*. Retrieved from <https://www.ncbex.org/statistics-and-research/covington-award>
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112, 155–159.
- Crosier, R. B. (1988). Multivariate generalizations of cumulative sum quality-control schemes. *Technometrics*, 30, 291–303.
- DeMars, C. E. (2004). Detection of item parameter drift over multiple test administrations. *Applied Measurement in Education*, 17, 265–300.
- Donoghue, J. R., & Isham, S. P. (1998). A comparison of procedures to detect item parameter drift. *Applied Psychological Measurement*, 22(1), 33–51.
- Goldstein, H. (1983). Measuring changes in educational attainment over time: Problems and possibilities. *Journal of Educational Measurement*, 20, 369–377.

- Guo, H., Robin, F., & Dorans, N. (2017). Detecting item drift in large-scale testing. *Journal of Educational Measurement*, 54, 265–284.
- Healy, J. D. (1987). A note on multivariate CUSUM procedures. *Technometrics*, 29, 409–412.
- Hotelling, H. (1931). The generalization of Student's ratio. *Annals of Mathematical Statistics*, 2, 360–378. doi: 10.1214/aoms/1177732979
- Huggins-Manley, A. C. (2017). Psychometric consequences of subpopulation item parameter drift. *Educational and Psychological Measurement*, 2017, 143–164. doi: 10.1177/0013164416643369
- Kang, H.-A., Zheng, Y., & Chang, H.-H. (2020). Online calibration of a joint model of item responses and response times in computerized adaptive testing. *Journal of Educational and Behavioral Statistics*, 45, 175–208.
- Klein Entink, R. H., Kuhn, J.-T., Hornke, L. F., & Fox, J.-P. (2009). Evaluating cognitive theory: A joint modeling approach using responses and response times. *Psychological Methods*, 14, 54–75.
- Lai, T. (1991). Asymptotic optimality of generalized sequential likelihood ratio tests in some classical sequential testing problems. In B. K. Ghosh & P. K. Sen (Eds.), *Handbook of sequential analysis* (pp. 121–144). New York: Marcel Dekker, Inc.
- Lee, Y.-H., & Lewis, C. (2021). Monitoring item performance with CUSUM statistics in continuous testing. *Journal of Educational and Behavioral Statistics*, 46, 611–648. doi: 10.3102/1076998621994563
- Liu, C., Han, K. T., & Li, J. (2019). Compromised item detection for computerized adaptive testing. *Front. Psychol.*, 10, 829. doi: 10.3389/fpsyg.2019.00829
- Lowry, C. A., Woodall, W. H., Champ, C. W., & Rigdon, S. E. (1992). A multivariate EWMA control chart. *Technometrics*, 34, 46–53.
- Marianti, S., Fox, J.-P., Avetisyan, M., Veldkamp, B. P., & TijmstraFirs, J. (2014). Testing for aberrant behavior in response time modeling. *Journal of Educational and*

- Behavioral Statistics*, 39, 426–451.
- Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41, 100–115. doi: 10.1093/biomet/41.1-2.100
- Pignatiello, J. J., & Runger, G. C. (1990). Comparisons of multivariate CUSUM charts. *Journal of Quality Technology*, 22, 173–186.
- Segall, D. O. (2002). An item response model for characterizing test compromise. *Journal of Educational and Behavioral Statistics*, 27, 163–179.
- Segall, D. O. (2004). A sharing item response theory model for computerized adaptive testing. *Journal of Educational and Behavioral Statistics*, 29, 439–460.
- Shu, Z., Henson, R., & Luecht, R. (2013). Using deterministic, gated item response theory model to detect test cheating due to item compromise. *Psychometrika*, 78, 481–497.
- Sinharay, S., & Johnson, M. S. (2020). The use of item scores and response times to detect examinees who may have benefited from item preknowledge. *British Journal of Mathematical and Statistical Psychology*, 73, 397–419.
- Tendeiro, J. N., Meijer, R. R., Schakel, L., & Maij-de Meij, A. M. (2013). Using cumulative sum statistics to detect inconsistencies in unproctored internet testing. *Educational and Psychological Measurement*, 73, 143–161.
- van der Linden, W. J. (2006). A lognormal model for response times on test items. *Journal of Educational and Behavioral Statistics*, 31, 181–204.
- van der Linden, W. J. (2007). A hierarchical framework for modeling speed and accuracy on test items. *Psychometrika*, 72, 287–308.
- van der Linden, W. J., & Guo, F. (2008). Bayesian procedures for identifying aberrant response-time patterns in adaptive testing. *Psychometrika*, 73(3), 365–384.
- van Krimpen-Stoop, E. M. L. A., & Meijer, R. R. (2001). CUSUM-based person-fit statistics for adaptive testing. *Journal of Educational and Behavioral Statistics*, 26, 199–218.
- Veerkamp, W. J. J., & Glas, C. A. W. (2000). Detection of known items in adaptive

- 1 testing with a statistical quality control method. *Journal of Educational and*
2 *Behavioral Statistics*, 25, 373–389.
- 3 Wang, X., & Liu, Y. (2020). Detecting compromised items using information from secure
4 items. *Journal of Educational and Behavioral Statistics*, 45, 667–689.
- 5 Wells, C. S., Subkoviak, M. J., & Serlin, R. C. (2002). The effect of item parameter drift
6 on examinee ability estimates. *Applied Psychological Measurement*, 26, 77–87.
- 7 Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing
8 composite hypotheses. *The Annals of Mathematical Statistics*, 9, 60–62.
- 9 Woodall, W. H., & Ncube, M. M. (1985). Multivariate CUSUM quality control procedures.
10 *Technometrics*, 27, 285–292.
- 11 Yang, Y., Ferdous, A., & Chin, T. Y. (2007, April). Exposed items detection in personnel
12 selection assessment: An exploration of new item statistic. Chicago, IL: Paper
13 presented at the annual meeting of the National Council of Measurement in
14 Education.
- 15 Zhang, J. (2014). A sequential procedure for detecting compromised items in the item pool
16 of CAT system. *Applied Psychological Measurement*, 38, 87–104.
- 17 Zhang, J., & Li, J. (2016). Monitoring items in real time to enhance CAT security. *Journal*
18 *of Educational Measurement*, 53, 131–151.
- 19 Zhang, J., Li, Z., & Wang, Z. (2010). A multivariate control chart for simultaneously
20 monitoring process mean and variability. *Computational Statistics and Data*
21 *Analysis*, 54, 2244–2252.
- 22 Zopluoglu, C. (2019). Detecting examinees with item preknowledge in large-scale testing
23 using extreme gradient boosting (xgboost). *Educational and Psychological*
24 *Measurement*, 79, 931–961.

Figure 2. Item Flagged as Drift by Joint Monitoring Procedures
 t : Evaluation points. Avg Abs Diff: Average of absolute differences between the re-estimated and initial item parameter values across the four parameters. Prop Correct: Proportion of correct responses in the calibration sample. Avg RT: Average of response times in the calibration sample.

Figure 3. Item Flagged as Drift by Response-based Monitoring Procedures
 t : Evaluation points. Avg Abs Diff: Average of absolute differences between the re-estimated and initial item parameter values across the four parameters. Prop Correct: Proportion of correct responses in the calibration sample. Avg RT: Average of response times in the calibration sample.