Setting the control limit at release for stability assurance

Gregory Chen, X.; van der Vaart, A.W.

DOI
10.1002/pst.2259

Publication date
2022

Document Version
Final published version

Published in
Pharmaceutical Statistics

Received: 4 January 2022 | Revised: 16 May 2022 | Accepted: 1 July 2022

DOI: 10.1002/pst.2259

MAIN PAPER

WILEY

Setting the control limit at release for stability assurance

X. Gregory Chen[1] | A. W. van der Vaart[2]

[1]Biostatistics and Research Decision Sciences, MSD, Zurich, Switzerland

[2]Delft Institute of Applied Mathematics, TU Delft, Delft, The Netherlands

Correspondence
X. Gregory Chen, Biostatistics and Research Decision Sciences, MSD, Zurich, Switzerland.
Email: stat1013@gmail.com and xiangyi.gregory.chen@msd.com

Funding information
European Research Council Grant, Grant/Award Number: 320637; The Netherlands Organisation for Scientific Research (NWO), Grant/Award Number: Spinoza Prize

Abstract

A common task in quality control is to determine a control limit for a product at the time of release that incorporates its risk of degradation over time. Such a limit for a given quality measurement will be based on empirical stability data, the intended shelf life of the product and the stability specification. The task is particularly important when the registered specifications for release and stability are equal. We discuss two relevant formulations and their implementations in both a frequentist and Bayesian framework. The first ensures that the risk of a batch failing the specification is comparable at release and at the end of shelf life. The second is to screen out batches at release time that are at high risk of failing the stability specification at the end of their shelf life. Although the second formulation seems more natural from a quality assurance perspective, it usually renders a control limit that is too stringent. In this paper we provide theoretical insight in this phenomenon, and introduce a heat-map visualisation that may help practitioners to assess the feasibility of implementing a limit under the second formulation. We also suggest a solution when infeasible. In addition, the current industrial benchmark is reviewed and contrasted to the two formulations. Computational algorithms for both formulations are laid out in detail, and illustrated on a dataset.

KEYWORDS

internal release limit, quality control, random coefficients model, release specification, stability specification

1 | INTRODUCTION

Critical quality attributes (CQAs) are physical, chemical, biological or microbiological properties of a drug that are critical to patient safety and drug efficacy (see ICH Q8-R2[1]). Determination of what attributes are CQAs depends on the nature of the drug, knowledge gained through development and commercialisation, and regulatory recommendation. Before a batch of drug products can be released to the market, these CQAs are measured on samples taken from the batch, and ascertained to meet certain criteria. These are described in a *release specification* (*RS*). Sampling and measurement are performed according to a registered laboratory procedure for every CQA. Since drugs may degrade over time, separate criteria ensure that a drug meets its quality specification until the end of its shelf life, referred to as the *stability specification* (***SS***). Failing to meet the SS can lead to recall of a batch from the market. If a CQA of a drug product is expected to change significantly over time, its release and stability specifications are usually different. Although in Japan and the US this may be true only for in-house criteria and the regulatory RS and SS are the same, in the European Union there is a regulatory requirement for distinct specifications for release and shelf-life.[2] A similar recommendation is given by the WHO.[3]

Typically, the specification for a CQA consists of an acceptable range for a measured sample (also called a *reportable value*) of a batch, denoted as y, such as $y \geqslant \gamma_{low}$, $y \leq \gamma_{upp}$ or $\gamma_{low} \leq y \leq \gamma_{upp}$, where γ_{low} and γ_{upp} are registered specifications. Besides the theoretical "safe zone" for a CQA, the specification need to take account of inevitable random variation in the manufacturing process and laboratory measurements.

In this paper we are concerned with determining a control limit for a reportable value at release (i.e., time zero) that incorporates potential degradation of the drug over time, given historical data, shelf life and stability specification (SS). We make our objectives explicit in the next section. The control limit may serve as a reference to determine the release specification for a stability-indicating CQA in the pre-commercial phase (e.g., based on data from a formal stability study), or could be used as an *internal release limit*[4] after commercialisation. The latter is useful in a situation that equal RS and SS were registered, but non-negligible trend is observed in a larger dataset that becomes available after commercialisation.

The paper is organised as follows. Two problem formulations are presented in Section 2. A widely used benchmark approach, its probabilistic objective and its drawbacks, are discussed in Section 3. The underlying population model in this benchmark approach is oversimplified. We hence proceed by first set up a more realistic population model in Section 4, Then, in Sections 5 and 6, we present the probabilistic establishment of the two problem formulation in such population model. Inference procedures based on a sample are provided in the same sections. These procedures are illustrated on two example datasets in Section 7, and a brief conclusion is drawn in Section 8. Proofs of propositions and theoretical justifications are collected in Appendix A.

2 | PROBLEM FORMULATION

Let Y_{it} be a reportable value (i.e., sample measurement) of a CQA for batch i at time t, where $i = 1, \cdots, n$ and n is the number of measured batches. We assume a decreasing trend of the CQA over time, and want to ensure a stability specification (SS) (i.e., $\geqslant \gamma$) at the end of shelf life $T > 0$ by setting a lower limit ($\geqslant \eta$) at release ($t = 0$). Letting $Y_{n+1,t}$ denote a future reportable value, we make this precise in one of the following two objectives:

1. Conditional on Trend/CoT: A future batch passes the specification $\geqslant \gamma$ at time T with at least the probability that it passes the limit $\geqslant \eta$ at time 0, that is, $P(Y_{n+1,T} \geqslant \gamma) \geqslant \Pr(Y_{n+1,0} \geqslant \eta)$.
2. Conditional on Individual/CoI: Given that it passes the limit $\geqslant \eta$ at time 0 a future batch passes the limit $\geqslant \gamma$ with high probability at time T, that is, $P(Y_{n+1,T} \geqslant \gamma \mid Y_{n+1,0} \geqslant \eta)$ is large.

It turns out to be an ambitious task to assure stability by controlling just a sample measurement at time 0. The key premises are (a) the random measurement error is not high, and (b) a future batch has a similar stability trend as the historical batches. The second premise is usually satisfied, because the drug batches are all manufactured and stored in the same manner. However, in modelling it is too strong to assume all batches follow exactly the same degradation pattern. More realistically, we expect that there is a typical pattern overall, while each individual batch still has a certain idiosyncrasy. In Section 4, we shall incorporate this in a mixed model.

The CoT formulation establishes a RS that is consistent with the SS in terms of riskiness. The riskiness of a batch with respect to the stability specification $\geqslant \gamma$ can be quantified by its failing rate $P(Y_{n+1,T} < \gamma)$, which depends on its true quality level $\mu_T = \mathrm{E}Y_{n+1,T}$ (batch mean) at time T and laboratory measurement error. It is usually plausible that the variance of the measurement error is constant over time, whence the riskiness of a batch changes only through its true quality level μ_t. If RS and SS would be the same and a non-negligible degradation exists, then a batch might be risky with respect to the SS but not to the RS. The choice of η in the CoT formulation prevents an increase in riskiness. Tightening the stability specification by the average stability degradation (e.g., setting $\eta = \gamma - bT$, based on an assumption of linear trend with slope b) can be viewed as a special case of the CoT formulation, see Section 5.1. Hence, although this probabilistic formulation does not sound intuitive, in practice one encounters it more often than expected.

The CoI formulation seems to give the most natural objective for quality assurance, which corresponds to filtering out a correct proportion of batches at time 0 that are at greatest risk of failing at time T. However, when $Y_{n+1,0}$ and $Y_{n+1,T}$ are only weakly correlated, and hence $Y_{n+1,0}$ carries not much information about $Y_{n+1,T}$, the CoI formulation becomes spurious. A weak correlation can be caused by large measurement error and/or large variation between the individual batches. If the trends of the empirical batches vary vastly, then the quality level in the future will be hard to

predict at time 0 even without measurement error, and the further the future, the less reliable the prediction. We discuss this further in Section 6.

In this study we consider a (linearly) decreasing trend and lower-bounded SS intervals. Generalisation to an increasing trend and upper-bounded SS intervals is obvious. In the unlikely situation that the trend is in the direction of improving CQA, there is no stability risk, and the problem of setting control limits becomes spurious.

3 | BENCHMARK APPROACH IN THE INDUSTRY

There is a scarce literature on controlling stability risk by setting a specification limit at time 0. The most prevalent approach seems to be due to Allen, Dukes and Gerger (henceforth referred to as the ADG approach).[6] They proposed, in the case that the trend in the CQA is linearly decreasing, to set the lower limit at time 0 of the interval $(\geqslant \eta_{ADG})$ equal to

$$\eta_{ADG} = \gamma - \widehat{b}T + t_{0.95,n^*}\sqrt{T^2 s^2_{\widehat{b}} + \frac{\widehat{\sigma}_e^2}{k}}, \tag{1}$$

where γ is the lower-sided SS, T is the shelf life (in months), $\widehat{b}$ is an estimate of the degradation slope across historical batches in the data, $s_{\widehat{b}}$ is the standard error of the estimator $\widehat{b}$, $\widehat{\sigma}_e^2/k$ is an estimate of laboratory method (assay) variance, k is the number of replicates to be averaged for a reportable value, and $t_{0.95,n^*}$ is the 95% percentile of the student t distribution with n^* degrees of freedom. Because the estimates $\widehat{b}$ and $\widehat{\sigma}_e$ may not come from the same fitted model and may not be based on the same dataset, n^* is calculated by the Satterthwaite approximation.

An intuitive interpretation of (1) is straightforward: increase the current SS $(\geqslant \gamma)$ by the amount of the average degradation during shelf life (note that $-\widehat{b}T > 0$), and further tighten the interval by an error margin reflecting the uncertainty in the estimated slope and the measurement error.

Because the original paper did not present a population model and objective function, a precise interpretation by an associated probabilistic statement is tenuous, but our best understanding of the procedure is as follows. Assume that the observation $Y_{n+1,T}$ of a future batch at time T satisfies $Y_{n+1,T} = a + bT + e$, where a is the unknown true batch mean at time 0, b is the batch slope of degeneration over time, and e is a residual error, assumed to possess a normal distribution $N\left(0, \sigma_e^2/k\right)$. Assume that $\widehat{b}$ is an unbiased estimator of b with a normal distribution with variance $\sigma^2_{\widehat{b}}$, and that $s_{\widehat{b}}$ and $\widehat{\sigma}_e$ are estimators of $\sigma_{\widehat{b}}$ and σ_e. Seeking the minimum value of a such that $P(Y_{n+1,T} \geqslant \gamma) \geqslant 0.95$, is equivalent to solving the following inequality for a:

$$P\left(\frac{Y_{n+1,T} - a - \widehat{b}T}{\sqrt{T^2 s^2_{\widehat{b}} + \frac{\widehat{\sigma}_e^2}{k}}} \geqslant \frac{\gamma - a - \widehat{b}T}{\sqrt{T^2 s^2_{\widehat{b}} + \frac{\widehat{\sigma}_e^2}{k}}}\right) = 0.95, \tag{2}$$

where $Y_{n+1,T} - a - \widehat{b}T = \left(b - \widehat{b}\right)T + e$ is normally distributed with mean 0 and variance $\sigma^2_{\widehat{b}} + \sigma_e^2/k$. By unbiasedness of $\widehat{b}$, the mean of $\left(b - \widehat{b}\right)T + e$ is zero, while its variance is $\sigma^2_{\widehat{b}}T^2 + \sigma_e^2$. Therefore, the statistic inside the probabilityshould (approximately) follow the Student t distribution with n^* degrees of freedom. Equation (2) is, therefore, solved for a, by setting the quotient on the right of the inequality inside the probability equal to the 0.95-quantile of this Student t distribution. In other words, the minimal value of a is the right side of (1).

Since a can be interpreted as the expectation at time zero, this reasoning suggests that the ADG approach sets a lower limit for the population mean of the reportable value at zero. As this population quantity is not observable, this may not be an attractive target. An alternative perspective on the method, which is in a similar spirit as the CoT formulation, is elaborated in the next section. Differences are that (1) the ADG approach does not explicitly model the batch-to-batch slope variation, thus underestimating estimation error as the number of empirical batches increases, and (2) the ADG approach lumps the variation of the lab measuring method and trend variation among batches in a single residual variance.

4 | OUR POPULATION MODEL

In the remainder of the paper we adopt the following linear random-coefficients model, which is realistic in our industrial setting.[4] A reportable value (a sample measurement) $Y_{i,t}$ from batch i at time t follows

$$Y_{i,t} = a + \alpha_i + (b + \beta_i)t + \epsilon_{i,t}. \tag{3}$$

The overall coefficients a and b are fixed effects, showing the average trend across batches. Each batch i has a random deviation in its individual intercept and trend introduced by

$$\alpha_i \sim N\left(0, \sigma_\alpha^2\right), \qquad \beta_i \sim N\left(0, \sigma_\beta^2\right).$$

The residual errors $\epsilon_{i,t} \sim N\left(0, \sigma_\epsilon^2\right)$ reflect measurement error. For simplicity we assume that all random effects $\alpha_i, \beta_j, \epsilon_{i,t}$ are mutually independent. The model is then parameterised by the vector $\theta := \left(a, b, \sigma_\alpha, \sigma_\beta, \sigma_\epsilon\right)$.

We make two remarks to justify this population model, one on linearity and the other on normality. (1) To detect a nonlinear pattern of degradation, we would need to measure the reported values on a fine time grid. However, in a formal stability study a batch is typically observed at a total of 7 time points, or at 4 or 5 time points in a follow-up stability study. This may not provide enough resolution to reliably distinguish a moderately nonlinear from a linear trend, let alone to distinguish between two types of nonlinear trends. Thus a linear trend might be the best we can handle in practice, unless some specific nonlinear trend is given strong scientific endorsement or is observed empirically during development. (2) Only few CQAs are known to follow non-Gaussian distributions. One of the most important examples is the relative potency for biologics (a ratio), which is known to follow a log-normal distribution. In such a case, one may seek an appropriate transformation of the reportable value, so that it follows the model given by Equation (3). Other CQAs with non-Gaussian distributions can be handled case by case, with appropriate modifications of the methodology proposed in this paper.

5 | COT FORMULATION

5.1 | CoT in population model

The objective of CoT is to ensure that the passing rate for $y \geqslant \gamma$ at a future time T is at least the passing rate for $y \geqslant \eta$ at the time of release. Under the population model (3), trend and intercept are random effects, and hence we cannot expect the objective to hold for every future batch. Instead we require it for at least a large proportion, say $100q\%$ of the future batches, where $q \in [0, 1]$ is left to the risk appetite in the specific setting.

In formula, the objective of CoT is then, given γ, to find η that fulfils

$$P_\theta\left[\frac{P_\theta\left(Y_{n+1,0} \geqslant \eta \mid \alpha_{n+1}, \beta_{n+1}\right)}{P_\theta\left(Y_{n+1,T} \geqslant \gamma \mid \alpha_{n+1}, \beta_{n+1}\right)} \leqslant 1\right] = q. \tag{4}$$

The probabilities in the quotient are conditional given the random effects, while the outer probability refers to the random effects. The parameter θ is treated as given. By some rearrangements (see Section 8), one can show that the batch intercept does not play a role in this inequality. In fact, the solution can be found explicitly and is given by

$$\eta_{CoT} = \gamma - \left[b + \Phi^{-1}(1-q)\sigma_\beta\right]T, \tag{5}$$

where $\Phi^{-1}(1-q)$ denotes the $(1-q)$th quantile of the standard normal distribution.

An intuitive interpretation of the formula is that it tightens the specification $\geqslant \gamma$ by the expected degradation up to time T assuming that the trend follows the q^{th} steepest slope from the batch slope distribution. In the trivial case when $\sigma_\beta \approx 0$, that is, all batches have almost the same trend, we find that $\eta \approx \gamma - bT$, which is the intuitive solution to balance the stringency of specifications at time 0 and T given linear degradation.

5.2 | Sample Inference for CoT

In formula (5), the quantities b and σ_β are unknown population parameters. The simplest solution would be to replace them by sample estimates $\widehat{b}$ and $\widehat{\sigma}_\beta$ given the historical data, but this would not take into account the uncertainty due to parameter estimation. Alternatively, in a Bayesian framework, a posterior distribution of η_{CoT} can be derived from the joint posterior distribution of (b, σ_β) in combination with formula (5). A suitable measure of location of that posterior, say the posterior mean or median, may be used as the final estimate. For a conservative solution, we may use the ξ^{th} posterior quantile for some $\xi > 0.5$, the choice $\xi = 0.5$ leading back to the posterior median. In practice, these quantities may be computed from a large set of sample values $\eta_{CoT}(b, \sigma_\beta)$ obtained by applying the map (5) to a sample from the posterior distribution of (b, σ_β).

6 | COI FORMULATION

6.1 | CoI in population model

The objective of CoI is to set a control limit $\geqslant \eta$ at time 0, so that the event $Y_{n+1,0} \geqslant \eta$ indicates a high probability that the event $Y_{n+1,T} \geqslant \gamma$ will occur at the later time point T. This can be expressed in a formula as the requirement on the conditional probability, for given γ:

$$P_\theta(Y_{n+1,T} \geqslant \gamma \mid Y_{n+1,0} \geqslant \eta_{CoI}) \geqslant q, \quad \eta_{CoI} \geqslant \gamma, \tag{6}$$

where $\theta = (a, b, \sigma_\alpha, \sigma_\beta, \sigma_\epsilon)$, and q is a prescribed, desired level of assurance, which is typically set close to 1. The constraint $\eta_{CoI} \geqslant \gamma$ ensures that the solution η_{CoI} is practically meaningful.

The function $\eta \mapsto P_\theta(Y_{i,T} \geqslant \gamma | Y_{i,0} \geqslant \eta)$ is increasing from $P_\theta(Y_{i,T} \geqslant \gamma)$ to (typically) 1 (see Lemma 6.2 below and Figure 1). Hence a solution η_{CoI} to the equation will exist for $q \geqslant P_\theta(Y_{i,T} \geqslant \gamma)$, and it will satisfy the restriction $\eta_{CoI} \geqslant \gamma$ if $q \geqslant P_\theta(Y_{i,T} \geqslant \gamma | Y_{i,0} \geqslant \gamma)$. For given η_{CoI}, a proportion $P_\theta(Y_{i,0} < \eta_{CoI})$ of produced batches will be rejected. As illustrated in Figure 1, this proportion may be large. We examine the practicality of the CoI formulation from this perspective in the next subsection.

We now derive the solution η_{CoI} to (6), assuming that all the parameters in (3) are known. Under model (3), the pair (Y_{i,t_1}, Y_{i,t_2}), for two given time values $t_1 < t_2$, follows a bivariate Gaussian distribution, given by

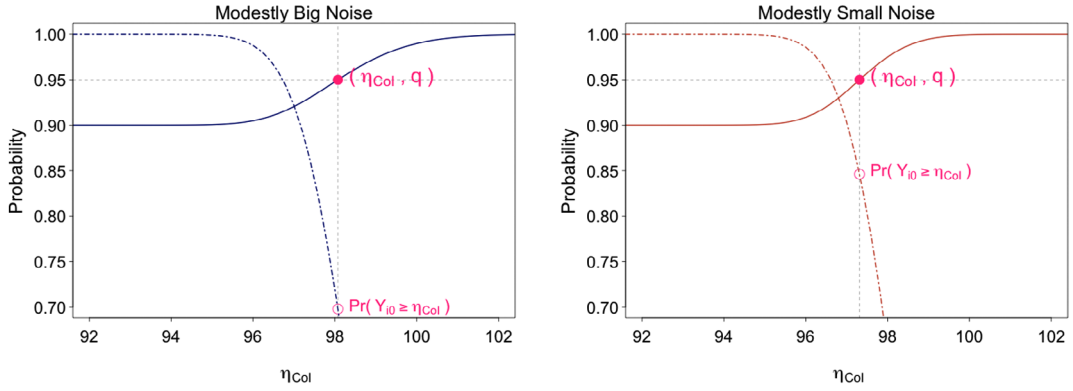

FIGURE 1 The curves $\eta_{CoI} \mapsto P_\theta(Y_{i,T} \geqslant \gamma | Y_{i,0} \geqslant \eta_{CoI})$ (solid) and $\eta_{CoI} \mapsto P_\theta(Y_{i,0} \geqslant \eta_{CoI})$ (dot dash), with η_{CoI} on the horizontal axis, for two different parameter settings and $\gamma = 95$. To obtain stability assurance at $q = 0.95$ the value of η_{CoI} must be set to the value at the vertical dotted line. The height of the dot dash curve on this line gives the proportion of accepted batches at time 0. Parameter settings: (Case I, Modestly Big Noise): $a = 98.69, b = -0.0635, \sigma_\alpha = 1, \sigma_\beta = 0.05, \sigma_\epsilon = 0.655$, which gives intra-correlation $\rho_{int} = \sigma_\alpha^2/(\sigma_\alpha^2 + \sigma_\epsilon^2) = 0.7$, $P(Y_{i,T} \geqslant \gamma) = 0.9$ and $P(Y_{i,0} \geqslant \gamma) = 0.999$. (Case II, Modestly Small Noise): $a = 98.45, b = -0.0729, \sigma_\alpha = 1, \sigma_\beta = 0.03, \sigma_\epsilon = 0.5$, which gives intra-correlation $\rho_{int} = 0.8$, $P(Y_{i,T} \geqslant \gamma) = 0.9$ and $P(Y_{i,0} \geqslant \gamma) = 0.999$; shelf life was $T = 24$ months in both cases.

$$\begin{pmatrix} Y_{i,t_1} \\ Y_{i,t_2} \end{pmatrix} \sim N\left(\begin{pmatrix} a+bt_1 \\ a+bt_2 \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2+\sigma_\beta^2 t_1^2+\sigma_\epsilon^2 & \sigma_\alpha^2+\sigma_\beta^2 t_1 t_2 \\ \sigma_\alpha^2+\sigma_\beta^2 t_1 t_2 & \sigma_\alpha^2+\sigma_\beta^2 t_2^2+\sigma_\epsilon^2 \end{pmatrix} \right). \tag{7}$$

Equation (6) is identical to, for $t_1 = 0$ and $t_2 = T$,

$$\frac{P_\theta(Y_{i,t_2} \geqslant \gamma, Y_{i,t_1} \geqslant \eta_{CoI})}{P_\theta(Y_{i,t_1} \geqslant \eta_{CoI})} = \frac{\int_{\eta_{CoI}}^{\infty}\int_{\gamma}^{\infty} \varphi_*(s_1,s_2)\,ds_2\,ds_1}{\int_{\eta_{CoI}}^{\infty} \varphi_{t_1}(s)\,ds} = q, \tag{8}$$

where φ_* is the density function of the bivariate Gaussian (7). In a different format,

$$\frac{\int_{\eta_{CoI}}^{\infty} P_\theta(Y_{i,t_2} \geqslant \gamma \,|\, Y_{i,t_1} = s)\,\varphi_{t_1}(s)\,ds}{P_\theta(Y_{i,t_1} \geqslant \eta_{CoI})} = q, \tag{9}$$

where φ_{t_1} is the marginal density function of $Y_{i,t_1} \sim N\left(a+bt_1, \sigma_\alpha^2+\sigma_\beta^2 t_1^2+\sigma_\epsilon^2\right)$. Equation (8) has no apparent analytical solution, but the terms in the quotient can be approximated numerically via an efficient algorithm. We use pmvnorm from the R package mvtnorm 1.0–11 for this task. As for Equation (9), it is known that the variate Y_{i,t_2} given $Y_{i,t_1} = s$ follows the $N(\mu_*, \sigma_*^2)$-distribution with

$$\mu_* = a + bt_2 + \frac{\sigma_\alpha^2+\sigma_\beta^2 t_1 t_2}{\sigma_\alpha^2+\sigma_\beta^2 t_1^2+\sigma_\epsilon^2}(s-a-bt_1), \tag{10}$$

$$\sigma_*^2 = \sigma_\alpha^2+\sigma_\beta^2 t_2^2+\sigma_\varepsilon^2 - \frac{\left(\sigma_\alpha^2+\sigma_\beta^2 t_1 t_2\right)^2}{\sigma_\alpha^2+\sigma_\beta^2 t_1^2+\sigma_\epsilon^2}. \tag{11}$$

Hence the left side of Equation (9) can be rewritten as $\int_\eta^\infty \left(1-\Phi\left(\frac{\gamma-\mu_*}{\sigma_*}\right)\right)\varphi_{t_1}(s)\,ds$, which involves an incomplete Gaussian integral which also has no analytical solution. Nevertheless, (9) is a useful expression, because $P_\theta(Y_{i,t_2} \geqslant \gamma \,|\, Y_{i,t_1} = s)$ is the answer to another common inquiry in quality assessment: given the release data of batch i, what is the chance for this batch to pass the SS at time T?

Wei[7] approached the problem in a similar way as (8) and called this the "unconditional rule." However, he formulated a different objective than (6) and focused on a 95% confidence limit of the predicted $Y_{i,T}$ (i.e., a future individual reportable value), instead of on controlling $Y_{i,T}$.

6.2 | Practicality of CoI

Unlike the CoT solution (5), the CoI solution has no clear analytical form that can be easily interpreted. Therefore we make some extra effort to analyse its characteristics. This also facilitates to understand when the solution might not be helpful in practice, even if its formulation is an honest translation of our interests.

We present a lemma and next three insights. A proof of the lemma can be found in the Appendix A.

Lemma 1. If (Y_1, Y_2) is a bivariate normal random vector with strictly positive correlation ρ, then:

i. the map $\eta \mapsto P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta)$ is continuous and increasing from $P(Y_2 \geqslant \gamma)$ at $\eta = -\infty$ to 1 at $\eta = \infty$, for any $\gamma \in \mathbb{R}$.
ii. the map $\eta \mapsto P(Y_2 \geqslant \gamma | Y_1 = \eta)$ is continuous and increasing from 0 at $\eta = -\infty$ to 1 at at $\eta = \infty$, for any $\gamma \in \mathbb{R}$.

iii. $P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) \geqslant P(Y_2 \geqslant \gamma | Y_1 = \eta)$, for any $\gamma, \eta \in \mathbb{R}$.

iv. both $P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta)$ and $P(Y_2 \geqslant \gamma | Y_1 = \eta)$ tend to $P(Y_2 \geqslant \gamma)$ as the correlation between Y_1 and Y_2 tends to 0, for any $\gamma, \eta \in \mathbb{R}$.

v. the solution η_ρ to $P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) = q$, for given γ and q with $q > P(Y_2 \geqslant \gamma)$, satisfies $(\eta_\rho - \mu_1)/\sigma_1 = \rho^{-1}(\Phi^{-1}(q) + (\gamma - \mu_2)/\sigma_2) + O(1)$, as $\rho \downarrow 0$, where $\mu_i = \mathrm{E}Y_i$ and $\sigma_i^2 = varY_i$. The same is true for the solution to $P(Y_2 \geqslant \gamma | Y_1 = \eta) = q$.

[Insight 1]: Under model (3), and hence (7) with $t_1 = 0$ and $t_2 = T$, the correlation between $Y_{i,0}$ and $Y_{i,T}$ is equal to $\rho(0,T) = \sigma_\alpha^2 / \sqrt{\left(\sigma_\alpha^2 + \sigma_\epsilon^2\right)\left(\sigma_\alpha^2 + \sigma_\beta^2 T^2 + \sigma_\epsilon^2\right)}$. It will approach 0 if any of, or a combination of, the following occurs: (1) predict a far future, that is, large T, (2) unpredictable individual batch trend, that is, large σ_β, (3) big random noise from measurement error, that is, large σ_ϵ relative to the process variation σ_α. In every of these cases setting a limit at time 0 to control stability at time T is nearly impossible, due to the weak relation between the variables $Y_{i,0}$ and $Y_{i,T}$, in view of (iv). By (v) the control limit tends to ∞ as $\rho(0,T) \downarrow 0$ at the rate $1/\rho(0,T)$ on the standard scale. This situation also holds in the trivial case that all batches have almost the same mean value, that is, when σ_α is (nearly) zero.

[Insight 2]: If the passing rate $P(Y_{i,T} \geqslant \gamma)$ at time T is higher than the desired assurance q, there is no stability risk. This is because $P(Y_{i,T} \geqslant \gamma | Y_{i,0} \geqslant \eta)$ is lower bounded by $P(Y_{i,T} \geqslant \gamma)$, for every η, by (ii) of the lemma. Furthermore, the minimal value permitted under the constraint $\eta \geqslant \gamma$ is $P(Y_{i,T} \geqslant \gamma | Y_{i,0} \geqslant \gamma)$ and hence we may check if the latter conditional probability exceeds q. If so, there is no stability risk, and we can set $\eta_{CoI} = \gamma$.

[Insight 3]: An alternative criterion, but in the spirit of CoI, is to find the minimum reportable value of a batch at time 0 that assures a passing rate of at least $100q\%$ at time T, that is,

$$\eta_{Alt} = \min\{\eta \geqslant \gamma : P_\theta(Y_{i,T} \geqslant \gamma | Y_{i,0} = \eta) \geqslant q\}. \tag{12}$$

By (ii) of the lemma, this limit solves the equation $P_\theta(Y_{i,T} \geqslant \gamma | Y_{i,0} = \eta_{Alt}) = q$ unless it is on the boundary ($\eta_{Alt} = \gamma$). Part (iii) of the lemma shows that the limit will be tighter than the solution of (6) (i.e., $\eta_{Alt} \geqslant \eta_{CoI}$), for every given q. Since the conditional distribution of $Y_{i,T}$ given $Y_{i,0} = \eta$ is normal with mean $\mu_{*0} = (a + bT) + \rho_{int}(\eta - a)$ and variance $\sigma_{*0}^2 = \sigma_\alpha^2 + \sigma_\beta^2 T^2 + \sigma_\epsilon^2 - \rho_{int}\sigma_\alpha^2$, for $\rho_{int} = \sigma_\alpha^2/\left(\sigma_\alpha^2 + \sigma_\epsilon^2\right)$, the solution to (12) is

$$\eta_{Alt} = \max\left(\frac{\gamma - bT - (1 - \rho_{int})a + \Phi^{-1}(q)\sigma_{*0}}{\rho_{int}}, \gamma\right).$$

This inflates quickly to ∞ as $\rho_{int} \downarrow 0$, in the same way as η_{CoI}, by (v) of the lemma.

Insights 1 and 2 are particularly relevant for implementation in practice. The former reveals that the CoI formulation may be spurious, while the latter exhibits situations where setting the control η_{CoI} is not necessary. One can make a preliminary check (e.g., one for $\rho(0,T)$ and one for $P(Y_{iT} \geqslant \gamma)$) to decide the suitability and necessity of calculating η_{CoI}.

We now present an example analysis to illustrate how it can be assessed when the CoI formulation may be helpful. A key consideration for implementation is the trade-off between the desired quality assurance, governed by q, and the business sustainability, reflected by $P(Y_{i,0} \geqslant \eta_{CoI})$. Low $P(Y_{i,0} \geqslant \eta_{CoI})$ means that, at the determined control limit η_{Col}, an unacceptably high number of the manufactured batches will be rejected. From a business point of view this control is disastrous, as the true quality of these rejected batches may be sufficiently high (at least during the beginning of the period). As discussed, the tight limit in the CoI formulation could be due to the low correlation between $Y_{i,0}$ and Y_{iT}, which makes controlling Y_{iT} via $Y_{i,0}$ practically impossible.

Consider a hypothetical drug with CQA known to degrade over time, and with a registered SS of $95 \leq Y_{iT} \leq 105$. Consider two different process averages ($a = 100$ or $a = 98.5$), two different degrees of degradation ($bT = -2$ or $bT = -3$), a process variation (i.e., variation of between-batch mean) fixed at $\sigma_\alpha = 0.5$, a batch slope variation $\sigma_\beta T$ varying over a grid $[0.3, 1]$, and measurement error σ_ε varying over $[0.1, 1.5]$. In practice these settings can come from prior knowledge of the product and the laboratory method during development. For each scenario, we derived η_{CoI} for $q = 0.95$ and $q = 0.99$, and calculated and visualised the stringency $P(Y_{i,0} \geqslant \eta_{CoI})$. Figure 2 gives a heat map of the stringency

for the best process average ($a = 100$, in the middle of the SS), with the four panels referring to the four possible combinations of bT and q. The same results for the less advantageous process average ($a = 98.5$) are given in Figure 3.

In practice, prior knowledge of the product and the laboratory method during development will yield appropriate values for the preceding parameters. We can then assess the suitability of the CoI approach by mapping the parameter values to the appropriate square in the heat maps. If this square is in a column of deep-pink colour, the measurement error is too high, and we cannot effectively control the risk of $Y_{i,T}$ via the reportable value $Y_{i,0}$. As a remedy, we may improve the precision of the laboratory measurement, so that the status shifts to a column that has at least some green cells. If the square is deep-pink, but in a column with a mix of green and deep-pink cells, we may consider shortening the shelf life of the drug.

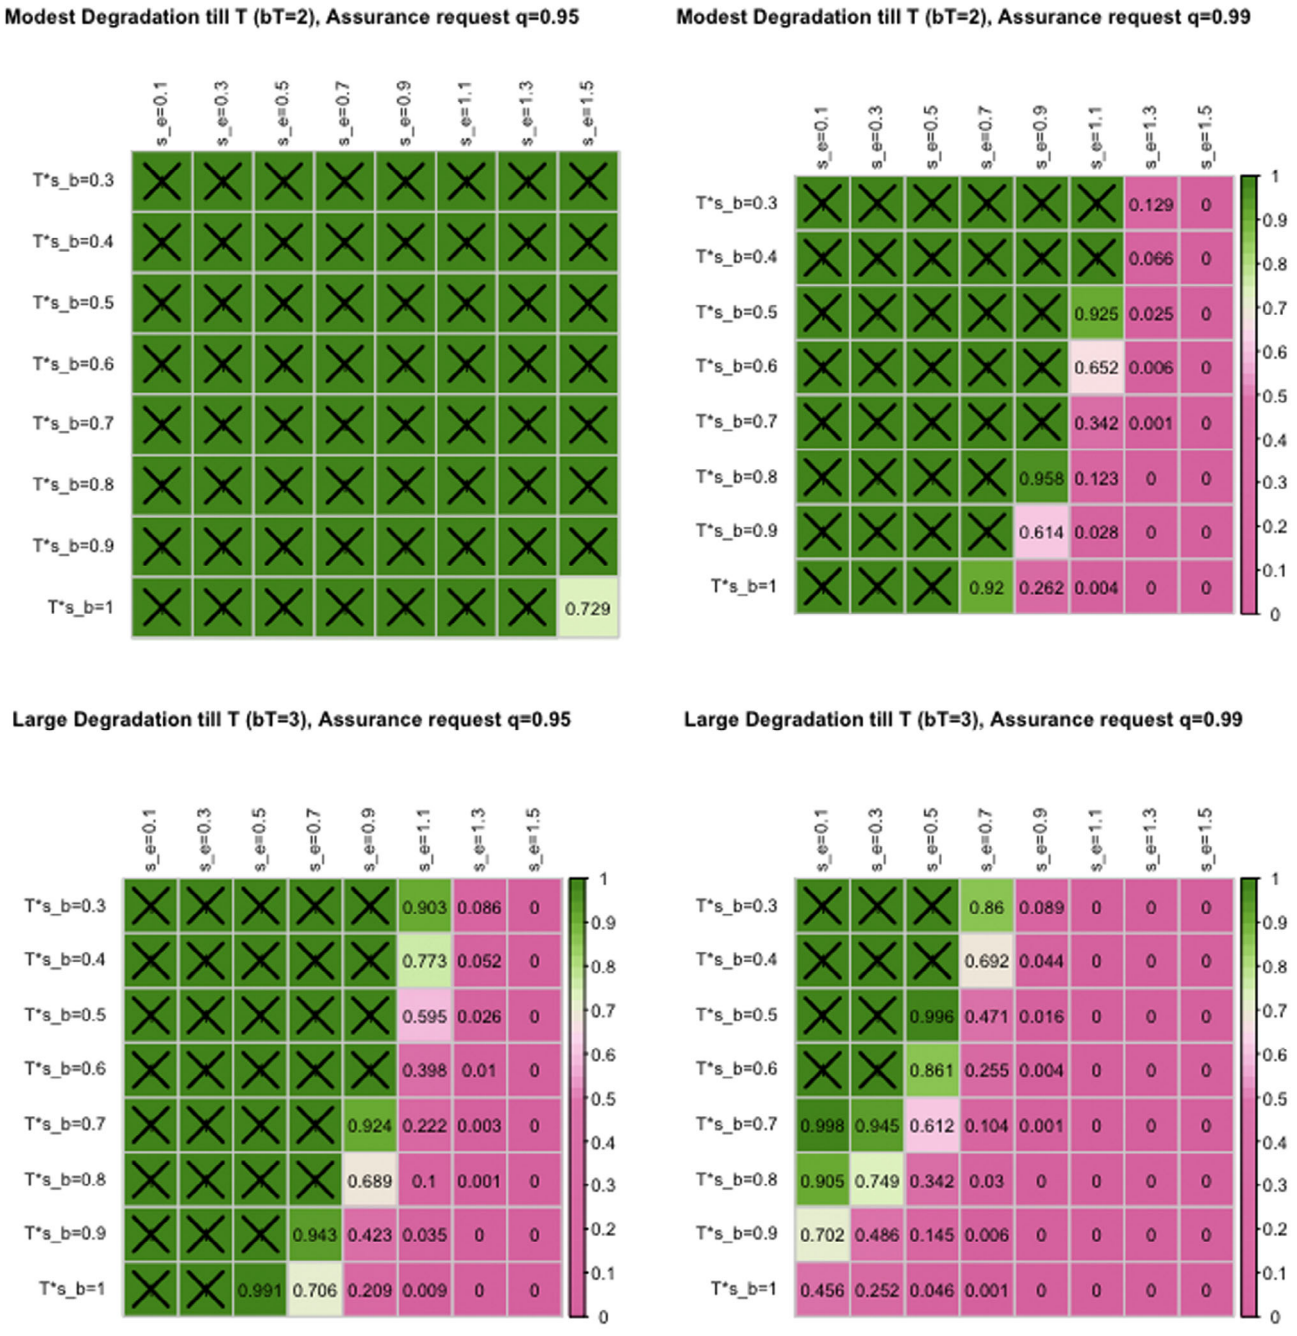

FIGURE 2 Stringency $P(Y_{i,0} \geqslant \eta_{CoI})$ indicated by colour scale under different scenarios. In all cases $a = 100$. The four panels concern the four combinations of the degradation ($bT = -2, -3$) and required quality assurance ($q = 0.95, 0.99$). Each panel gives combinations of $T\sigma_\beta$ (indicated by T*s_b) and σ_e (indicated by s_e). Cells marked by a cross are situations in which $P(Y_{iT} \geqslant \gamma) \geqslant q$ and hence no η_{CoI} needs to be calculated.

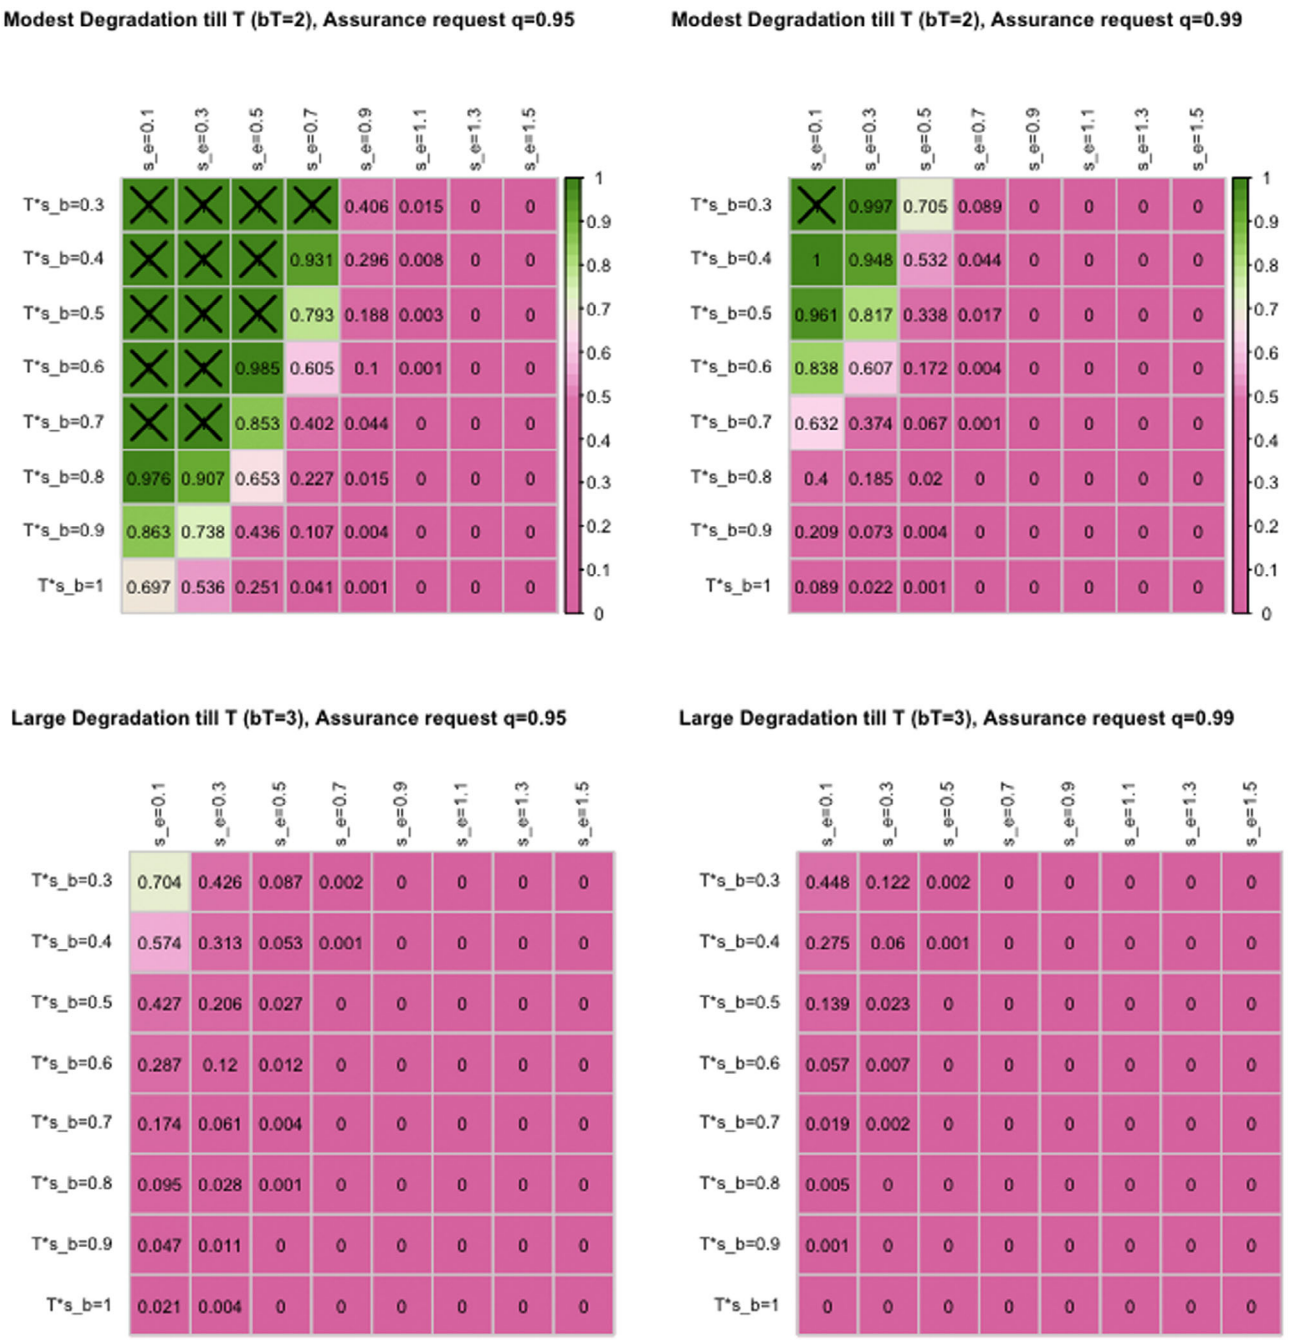

FIGURE 3 Stringency $P(Y_{i,0} \geqslant \eta_{CoI})$ indicated by colour scale under different scenarios. In all cases $a = 98.5$. The four panels concern the four combinations of the degradation ($bT = -2, -3$) and required quality assurance ($q = 0.95, 0.99$). Each panel gives combinations of $T\sigma_\beta$ (indicated by T*s_b) and σ_e (indicated by s_e). Cells marked by a cross are situations in which $P(Y_{iT} \geqslant \gamma) \geqslant q$ and hence no η_{CoI} needs to be calculated

In general, we find that the CoI approach quickly becomes impractical when the noise is large relative to the process variation σ_α, or the process average a is small, as indicated by the abundance of deep-pink squares.

6.3 | Sample inference for CoI

In practice, the population parameters θ in the model (3) are unknown, and must be estimated from sample data. To simplify the notation onwards, we fix $t_1 = 0$ and $t_2 = T$, where T is the shelf life in months (typically

$T \in \{18, 19, \cdots, 60\}$). We denote the empirical data by $\mathbf{y} = \{Y_{it} : i = 1, 2, \cdots, n;\ t = 0, \cdots, T\}$, indexed by batch ID i and observed time t.

From a frequentist perspective, we may utilise an estimator $\widehat{\theta}(\mathbf{y})$ of θ, and seek a (minimal) η_{CoI} that fulfils

$$\mathrm{E}_\theta\left[P_{\widehat{\theta}(\mathbf{y})}(Y_{n+1,T} \geqslant \gamma \mid Y_{n+1,0} \geqslant \eta_{CoI})\right] \geqslant q, \quad \text{for all } \theta.$$

The expectation pertains to the sample distribution of $\mathbf{y}$, which appears in $\widehat{\theta}(\mathbf{y})$. Because it seems complicated to solve such a set of equations exactly, we may determine an approximate solution by plugging in an estimator $\widehat{\theta}(\mathbf{y})$ of θ into the solution $\eta_{CoI}(\theta)$ of Equation (6), for given θ. Thus the control limit becomes $\eta_{CoI}\left(\widehat{\theta}(\mathbf{y})\right)$.

From a Bayesian perspective, there are two reasonable solutions, both based on the posterior distribution of θ given $\mathbf{y}$. The first, denoted **B1**, is to seek a (minimal) η_{CoI} that fulfils

$$\mathrm{E}[P_\theta(Y_{n+1,T} \geqslant \gamma \mid Y_{n+1,0} \geqslant \eta_{\mathrm{CoI}}) \mid \mathbf{y}] \geqslant q.$$

Here the expectation refers to the posterior distribution of θ given $\mathbf{y}$. An alternative Bayesian solution, denoted **B2**, is a measure of location of the posterior distribution of $\eta_{CoI}(\theta)$ given $\mathbf{y}$, which is induced by the posterior distribution of θ given $\mathbf{y}$ and the solution map $\theta \mapsto \eta_{\mathrm{CoI}}(\theta)$ of (6).

For computational purposes in both cases the posterior distribution of θ given $\mathbf{y}$ can be approximated by the empirical distribution of a sample $\{\theta_k\}_{k=1}^B$ of values obtained via an MCMC algorithm. The approximate solution **B1** is then given by

$$\min\left\{\eta : \frac{1}{B}\sum_{k=1}^{B} P_{\theta_k}(Y_{n+1,T} \geqslant \gamma \mid Y_{n+1,0} \geqslant \eta) \geqslant q\right\}. \tag{13}$$

The posterior distribution of $\eta_{CoI}(\theta)$ can be approximated by the empirical distribution of the values $\{\eta_{CoI}(\theta_k)\}_{k=1}^B$, and hence the **B2** solution by the location (e.g., median or quantile) of these values.

The frequentist and two Bayesian solutions are all different, but for large sample sizes will be similar, under mild conditions. The large sample properties of the three procedures are discussed in Section 8, where it is shown that the estimated control limits differ by no more than a centred normal variable with dispersion equal to the inverse sample size. For a conservative approach, the estimated control limit can be heightened by a quantile of the approximating normal distribution, thus yielding the limit $\eta_{CoI}\left(\widehat{\theta}(\mathbf{y})\right) + \widehat{\tau}_0\Phi^{-1}(\xi)/\sqrt{n}$, in the notation of Lemma 8. The Bayesian approach **B2** that uses the ξth quantile of the posterior distribution of η_{CoI} with $\xi > 0.5$ automatically incorporates a tightening of the control limit. Lemma 8 shows that this is asymptotically equivalent to the tightening of the frequentist procedure.

7 | CALCULATION WITH EXAMPLE DATASETS

We generated two datasets under the observational model (3), using two scenarios of parameter settings, and calculated η via the different approaches. For each scenario we set the shelf life T to 36 months, and SS equal to $\geqslant 95$, that is, $r = 95$. We generated 10 batches of data ($i = 1, ..., 10$), each with one observation at $t = 0, 6, 12, 18, 24$ months.

The parameters governing (3) are a, b, σ_α, σ_β and σ_e. We set $\rho = \sigma_\alpha^2/(\sigma_\alpha^2 + \sigma_\varepsilon^2)$, and write $\rho(0, T)$ for the correlation between $Y_{i,0}$ and $Y_{i,T}$. The two parameter settings were:

- (Modest intra-correlation) $P(Y_{i,0} \geqslant 95) = 0.9999$, $P(Y_{i,T} \geqslant 95) = 0.9$, $\sigma_\alpha = 0.5$, $\sigma_\beta = 0.01$, $\rho = 0.8$. This results in $a = 97.1$, $b = -0.034$, $\sigma_\epsilon = 0.25$ and a modest correlation $\rho(0, T) = 0.67$.
- (Weak intra-correlation) $P(Y_{i,0} \geqslant 95) = 0.9999$, $P(Y_{i,T} \geqslant 95) = 0.9$, $\sigma_\alpha = 0.5$, $\sigma_\beta = 0.04$, $\rho = 0.4$. This results in $a = 100.9$, $b = -0.068$, $\sigma_\epsilon = 0.6$ and a weak correlation $\rho(0, T) = 0.19$.

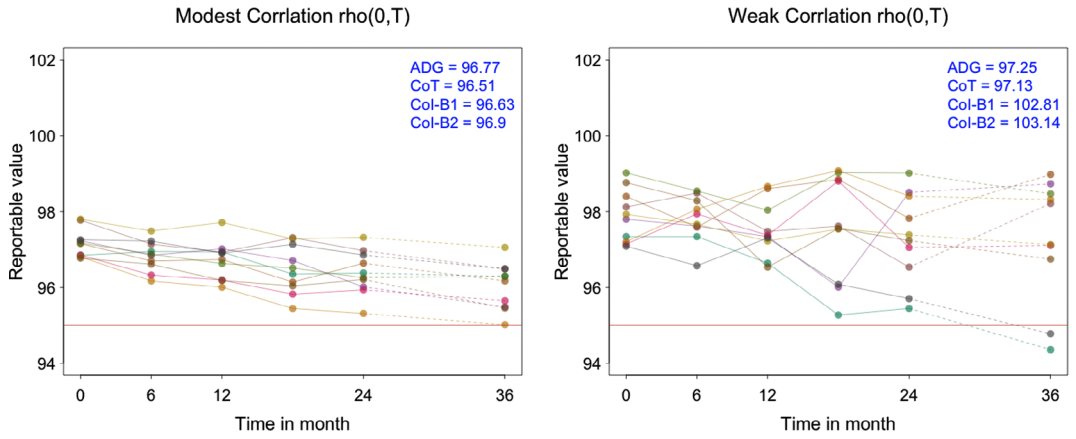

FIGURE 4 Graphical representation of data simulated using a scenario of modest ($\rho(0,T)=0.67$) and a scenario of weak ($\rho(0,T)=0.19$) correlation between reportable values at release and end of shelf life. Shown are also the control limits η at time 0 for four methods, for stability specification set at $\geqslant\gamma=95$. Data at 36 months are plotted for illustration, but were not included in the calculation of the control limits

For each dataset, η was calculated by four methods:

1. The ADG approach (1) with $q=0.95$. The estimator $\widehat{b}$ was taken equal to the ordinary least square estimator based on the $10\times 5=50$ observations Y_{it} with corresponding standard error s_b, ignoring the mixed model structure of the simulation scheme, using the function *lm* from R basic. The error variance σ_e was estimated as the sample variance of the regression residuals after correcting for a linear trend per batch. The value of k was set to 1.
2. The CoT solution as in (5) with plug-in estimators. We estimated the parameters b and σ_β by fitting a linear mixed model, using the function lmer from the R package {lme4}.[8] (Alternatively, Bayesian point estimators may be derived using the package {blme}.[9]) The parameter q was set to $q=0.8$, where it is noted that this parameter has a different meaning than in the other two approaches. In the CoT formulation q indicates the proportion of future batches for which the passing rate at time T is at least the passing rate at release, whereas in the other approaches q refers directly to the passing rate of an individual product. Setting it to 0.95 would lead to a very conservative *CoT* solution when the batch slope variation is not small. In practice, the values of q will be set case by case.
3. The two Bayesian solutions **B1** and **B2** for the CoI formulation, described in Section 6.3, with $q=0.95$. The posterior distribution of the parameters was approximated via an MCMC algorithm implemented in R package {brms},[10,11] with 2 chains, 2000 burn-in samples followed by $B=3000$ samples for each chain, using a $N(100,30^2)$ prior for a, a $N(0,5^2)$ prior for b, and an half-Cauchy prior[12] with scale 0.1 for σ_α, σ_β and σ_ϵ. The convergence and quality of the chain was checked graphically based on a trace plot, an ACF plot, and a Geweke diagnostic plot via R package {coda}[13] and the `launch_shinystan` function in {brms}.[14] The chosen priors are weakly informative. Experiments with other prior settings (e.g., variance 100^2 in the normal priors, and scale 0.01 in the half-Cauchy, which renders the priors even less informative), did not show significant changes in the posterior means, and are not reported. In practice, one may form weakly informative priors based on other empirical data or scientific expectation of model parameters. For CoI-**B1**, we used the average across the MCMC sample in (13). For CoI-**B2**, Equation (6) had no solution for some MCMC iterates θ_k (see Insight 2 in section 6.2), those $\eta_{Col}(\theta_k)$ were treated as some synthetic value <95, and the final estimate was taken to be the posterior median of $\{\eta_{Col}(\theta_k)\}_k$.

The example datasets and the calculated results are displayed in Figure 4. Data at 36 months are plotted for illustration, but were not included in the calculation of the control limits. One can see that with the dataset showing modest $\rho(0,T)$, the three approaches result in implementable control limits at time 0. However, when $\rho(0,T)$ is small, the CoI limits become impractically high.

8 | CONCLUSION

Adopting a realistic population model for a reportable value of a drug, we have discussed two formulations for setting a control limit at time 0 that incorporates stability risk. The two formulations correspond to different probabilistic objectives, which we worked out in a realistic random-effects model setting. Both formulation intend to determine a limit at release for an individual reportable value (of a CQA of a drug product). This is in line with how the limit will be applied in practice, and hence can be viewed as an improvement over the ADG approach, which determines a limit for an unobservable batch mean under a simplified population model, see Section 3.

The objective **CoT** seems to be generally appropriate. It determines a release specification that tightens the desired stability specification by a conservative estimate of average degradation over the course of shelf life. The level of conservativeness is controlled by choosing an appropriate quantile from the estimated distribution of the batch slope in a linear random-coefficients model. The **CoT** formula does not explicitly include the variability of individual measurements. Its use as a limit for an individual reportable value is justified under the assumption that the stability specification has already accounted for that source of variability. This assumption generally holds, since in practice the stability specification is also applied to an individual reportable value. (One should of course check that the reportable value at the stability test is consistent with the reportable value at the release test.)

The objective CoI is a natural translation of the aim of quality assurance, and focuses on filtering the batches at release on their reportable value. However, this aim becomes unattainable when the correlation $\rho(0, T)$ between the reportable values at time 0 and time T is near zero. The control limit at release may then be so tight that the majority of the batches will be rejected, as illustrated in Section 7.

We presented a visualisation by heat maps in Section 6.2 to help assess when the CoI approach will be helpful. The essence is to assess if the signal-to-noise ratio in the data is sufficiently high. The heat maps can be navigated to infer how to enhance the signal-to-noise ratio, by taking actions that lift the current status from a deep-pink cell to a green cell (e.g., improve the precision of the laboratory measurement to a certain level, or shorten the shelf life of the product by a certain number of months).

AUTHOR CONTRIBUTIONS

Both authors contributed to the full paper.

ACKNOWLEDGEMENTS

The research leading to these results was partly financed by European Research Council Grant Agreement 320637 and by an NWO Spinoza prize by the Netherlands Organisation for Scientific Research (NWO). The authors would also like to thank the associate editor and two anonymous reviewers for their constructive comments and suggestions.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

DATA AVAILABILITY STATEMENT

Only simulated data are used in this paper. Source code associated with this paper is publicly available on Github https://github.com/hoppanda/IRL.

ORCID

X. Gregory Chen https://orcid.org/0000-0003-0962-2466
A. W. van der Vaart https://orcid.org/0000-0002-8074-2375

REFERENCES

1. Holm P, Allesø M, Bryder MC, Holm R. Q8 (R2) pharmaceutical development. *ICH Qual Guidel: Implement Guide*. 2017;535-577.
2. Elder D. ICH Q6A specifications: test procedures and acceptance criteria for new drug substances and new drug products: chemical substances. In *ICH Qual Guidel: Implement Guide*. 2017;433-466. https://onlinelibrary.wiley.com/doi/10.1002/9781118971147.ch20
3. World Health Organization. Stability testing of active pharmaceutical ingredients and finished pharmaceutical products. *WHO Tech Rep Ser*. 2009;953:87-123.
4. Altan S, Dong Y, Gorko MA, et al. Understanding internal release limits. *Pharm Freeze Dry Technol*. 2018;42:20.
5. Avd V. Asymptotic statistics. *3 of Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press; 1998.

6. Allen PV, Dukes GR, Gerger ME. Determination of release limits: a general methodology. *Pharm Res*. 1991;8(9):1210-1213.
7. Wei GC. Simple methods for determination of the release limits for drug products. *J Biopharm Stat*. 1998;8(1):103-114.
8. Bates D, Mächler M, Bolker B, Walker S. Fitting linear mixed-effects models using lme4. *J Stat Softw*. 2015;67(1):1-48. doi:10.18637/jss.v067.i01
9. Dorie V, Dorie MV. Package 'blme'. Bayesian Linear Mixed-Effects Models. https://CRAN.R-project.org/package=blme2015.
10. Bürkner PC. Brms: an R package for Bayesian multilevel models using Stan. *J Stat Softw*. 2017;80(1):1-28. doi:10.18637/jss.v080.i01
11. Stan Development Team. RStan: The R Interface to Stan; 2022. R package version 2.21.5.
12. Gelman A. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal*. 2006;1(3):515-534.
13. Plummer M, Best N, Cowles K, Vines K. CODA: convergence diagnosis and output analysis for MCMC. *R News*. 2006;6(1):7-11.
14. Stan Development Team. Shinystan: Interactive Visual and Numerical Diagnostics and Posterior Analysis for Bayesian Models; 2017. R package version 2.4.0.

How to cite this article: Chen XG, van der Vaart AW. Setting the control limit at release for stability assurance. *Pharmaceutical Statistics*. 2022;1-19. doi:10.1002/pst.2259

APPENDIX A

A.1 | Proofs

Proof of formula (5). Under model (3), with $a^* = a + \alpha$ (drop the batch index i for brevity), Equation (4) can be rewritten as

$$P_{\alpha,\beta}\left[1-\Phi\left(\frac{\eta-a^*}{\sigma_e}\right) \leqslant 1-\Phi\left(\frac{\gamma-a^*-(b+\beta)t}{\sigma_e}\right)\right] = q$$
$$P_{\alpha,\beta}\left[\Phi\left(\frac{\gamma-a^*-(b+\beta)t}{\sigma_e}\right) \leqslant \Phi\left(\frac{\eta-a^*}{\sigma_e}\right)\right] = q$$
$$P_{\alpha,\beta}\left[\frac{\gamma-a^*-(b+\beta)t}{\sigma_e} \leqslant \frac{\eta-a^*}{\sigma_e}\right] = q$$

By some more algebra this can be rewritten further in the form

$$P_{\beta}\left[\frac{\gamma-\eta-bt}{\sigma_\beta t} \leqslant \frac{\beta}{\sigma_\beta}\right] = q.$$

Because β is normally distributed by assumption, it follows that $\gamma-\eta-bt = \Phi^{-1}(1-q)\sigma_\beta t$, which is equivalent to Equation (5).

Proof of Lemma 1. Because $P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) = P(Z_2 \geqslant (\gamma-\mu_2)/\sigma_2 | Z_1 \geqslant (\eta-\mu_1)/\sigma_1)$, for $Z_i = (Y_i-\mu_i)/\sigma_i$, and the same when conditioning on $Y_1 = \eta$, the situation can be reduced to the case that Y_1 and Y_2 possess mean zero and variance 1. Hence assume

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right].$$

Then (Y_1, Y_2) has density $f(y_1, y_2) \propto \exp\left\{-\frac{y_1^2-2\rho y_1 y_2+y_2^2}{2(1-\rho^2)}\right\}$, and the density of Y_2 given $Y_1 \geqslant \eta$ is given by

$$f_\eta(y) = \frac{\int_\eta^\infty f(s,y)\,ds}{P(Y_1 \geqslant \eta)} \propto \exp\left\{-\frac{y^2}{2(1-\rho^2)}\right\} \int_\eta^\infty \exp\left(\frac{\rho s y}{1-\rho^2}\right) d\mu(s)$$

for $d\mu(s) = \exp\left\{-\frac{s^2}{2(1-\rho^2)}\right\} ds$. Hence

$$\frac{f_\eta(y)}{f_{\eta'}(y)} \propto \frac{\int_\eta^\infty \exp(\tau s y)\,d\mu(s)}{\int_{\eta'}^\infty \exp(\tau s y)\,d\mu(s)}, \qquad \tau = \frac{\rho}{1-\rho^2}.$$

We first show that for $\eta > \eta'$, this function is increasing in y, or equivalently that its logarithm is increasing. The derivative of the latter function is

$$\frac{d}{dy}\log\frac{f_\eta(y)}{f_{\eta'}(y)} = \frac{\int_\eta^\infty \tau s \exp(\tau s y)\,d\mu(s)}{\int_\eta^\infty \exp(\tau s y)\,d\mu(s)} - \frac{\int_{\eta'}^\infty \tau s \exp(\tau s y)\,d\mu(s)}{\int_{\eta'}^\infty \exp(\tau s y)\,d\mu(s)} = \frac{\tau \int_{\eta'}^\infty \exp(\tau s y)\,d\mu(s)}{\int_\eta^\infty \exp(\tau s y)\,d\mu(s)} \left[\mathrm{E}\mathbf{1}_{s \geqslant \eta} S - \mathrm{E}\mathbf{1}_{s \geqslant \eta}\mathrm{E}S\right]$$

for S a random variable with density relative to $d\mu(s)$ given by

$$p(s) = \frac{\exp(\tau s y)\mathbf{1}_{s \geqslant \eta'}}{\int_{\eta'}^\infty \exp(\tau s y)\,d\mu(s)}.$$

The term in square bracket is $\mathrm{cov}\left(\mathbf{1}_{s \geqslant \eta}, S\right)$ and is non-negative, as $s \mapsto \mathbf{1}_{s \geqslant \eta}$ and $s \mapsto s$ are both increasing functions (see Lemma 8 below).

Now, for $Y \sim f_{\eta'}$.

$$P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) = \int_\gamma^\infty f_\eta(y)\,dy = \mathrm{E}_{\eta'}\mathbf{1}_{Y \geqslant \gamma}\frac{f_\eta(Y)}{f_{\eta'}(Y)} \geqslant \mathrm{E}_{\eta'}\mathbf{1}_{y \geqslant \gamma}\mathrm{E}_{\eta'}\frac{f_\eta(y)}{f_{\eta'}(y)} = P_{\eta'}(Y \geqslant \gamma)\cdot 1,$$

again by the covariance inequality. This finishes the proof of assertion (i).

Assertion (ii) is immediate from the form of the conditional distribution $Y_2 \mid Y_1 = \eta \sim N(\rho\eta, 1-\rho^2)$.

Assertion (iii) follows from

$$P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) = \frac{\int_\eta^\infty P(Y_2 \geqslant \gamma | Y_1 = s)\phi(s)\,ds}{P(Y_1 \geqslant \eta)} \geqslant \frac{\int_\eta^\infty P(Y_2 \geqslant \gamma | Y_1 = \eta)\phi(s)\,ds}{P(Y_1 \geqslant \eta)} = P(Y_2 \geqslant \gamma | Y_1 = \eta),$$

where the inequality follows from (i).

Assertion (iv) is a consequence of the fact that the correlation coefficient completely determines the dependence between multivariate Gaussian variables.

The assumption $P(Y_2 \geqslant \gamma) < q$ in the first part of (v) is equivalent to $\Phi^{-1}(q) + \gamma > 0$. If $0 < r < \Phi^{-1}(q) + \gamma$, then

$$\int_{r/(2\rho)}^\infty \Phi\left(\frac{s\rho - \gamma}{\sqrt{1-\rho^2}}\right)\phi(s)\,ds \leqslant \int_{r/(2\rho)}^{r/\rho} \Phi\left(\frac{r-\gamma}{\sqrt{1-\rho^2}}\right)\phi(s)\,ds + 1 - \Phi(r/\rho).$$

We conclude that, as $\rho \downarrow 0$,

$$P(Y_2 \geqslant \gamma | Y_1 \geqslant r/(2\rho)) \leqslant \Phi\left(\frac{r-\gamma}{\sqrt{1-\rho^2}}\right) + \frac{1-\Phi(r/\rho)}{1-\Phi(r/(2\rho))} \to \Phi(r-\gamma) < q,$$

where the convergence of the second term to zero follows, because by Mills ratio, as $\rho \downarrow 0$,

$$\frac{1-\Phi(r/\rho)}{1-\Phi(r/(2\rho))} \sim \frac{\phi(r/\rho)}{\phi(r/(2\rho))} \frac{r/(2\rho)}{r/\rho} \leqslant e^{-3r^2/(8\rho^2)} \to 0.$$

In view of the monotonicity (i), we conclude that $\eta_\rho \geqslant r/(2\rho)$, for sufficiently small ρ.

By similar reasoning, for any $\eta < \overline{\eta}$,

$$\Phi\left(\frac{\eta\rho-\gamma}{\sqrt{1-\rho^2}}\right) \leqslant P(Y_2 \geqslant \gamma | Y_1 \geqslant \eta) \leqslant \Phi\left(\frac{\overline{\eta}\rho-\gamma}{\sqrt{1-\rho^2}}\right) + \frac{1-\Phi(\overline{\eta})}{1-\Phi(\eta)}.$$

For $\eta = \eta_\rho$ and $\overline{\eta} = \eta_\rho + \sqrt{\rho}$, another application of Mills ratio shows that the ratio on the right is bounded above by $e^{-\eta_\rho\sqrt{\rho}} \leqslant e^{-r/\left(2\sqrt{\rho}\right)} \ll \rho$. We conclude that

$$\Phi\left(\frac{\eta_\rho\rho-\gamma}{\sqrt{1-\rho^2}}\right) \leqslant q, \qquad \Phi\left(\frac{\eta_\rho\rho+\rho^{3/2}-\gamma}{\sqrt{1-\rho^2}}\right) \geqslant q-\rho.$$

Solving this for η_ρ gives the first assertion of (v).

The solution to $P(Y_2 \geqslant \gamma | Y_1 = \eta) = q$ can be derived analytically, as $\eta = \left(\Phi^{-1}(q)\sqrt{1-\rho^2}+\gamma\right)/\rho$. This readily gives the second part of assertion of (v).

Lemma 2. If $f, g : \mathbb{R} \mapsto \mathbb{R}$ are non-decreasing functions, then $\operatorname{cov}(f(U), g(U)) \geqslant 0$, for any random variable U.

Proof. By monotonicity $(f(u)-f(v))(g(u)-g(v)) \geqslant 0$, for any u, v. Applying this to independent copies U, V of U, we find $\mathrm{E}(f(U)-f(V))(g(U)-g(V)) \geqslant 0$. The left side can be worked out as $2\mathrm{E}f(U)g(U) - 2\mathrm{E}f(U)\mathrm{E}g(U) = 2\operatorname{cov}(f(U), g(U))$.

A.2 | Asymptotics of the Bayesian estimators for CoI formulation

Let $\Theta \subset \mathbb{R}^d$ be open, and let $H : \Theta \times \mathbb{R} \to \mathbb{R}$ be a continuously differentiable map such that $\eta \mapsto H(\theta, \eta)$ is strictly increasing from a negative value to a positive value, for every $\theta \in \Theta$. Let $\eta : \Theta \to \mathbb{R}$ be defined by

$$H(\theta, \eta(\theta)) = 0.$$

We are interested in the plug-in estimator $\eta\left(\widehat{\theta}\right)$, given an estimator $\widehat{\theta}$ of θ; in the induced posterior distribution of $\eta(\theta)$, given a posterior distribution for θ; and in the solution to $\int H(\theta, \eta)\Pi(d\theta | Y_n) = 0$, for $\Pi(\theta \in \cdot\,|Y)$ a posterior distribution.

The intended application is the map $H(\theta, \eta) = \Pr_\theta(Y_T \geqslant \gamma | Y_0 \geqslant \eta) - q$, for a bivariate normal vector (Y_0, Y_T) with positive correlation, and numbers γ and q such that $\Pr(Y_T \geqslant \gamma) < q < 1$.

Write $H_\theta(\theta, \eta)$ for the $(d \times 1)$ vector of partial derivatives of H with respect to θ and $H_\eta(\theta, \eta)$ for the partial derivative with respect to η. We set $\eta_0 = \eta(\theta_0)$ and assume throughout that $H_\eta(\theta_0, \eta_0) > 0$.

Let $\widehat{\theta}_n$ be estimators based on data Y_n, for $n = 1, 2....$ The following lemma shows that asymptotic normality of these estimators carries over onto the solutions $\eta\left(\widehat{\theta}_n\right)$. Write $\rightsquigarrow$ for convergence in distribution of random vectors. In particular, $Z_n \rightsquigarrow N(0, \Sigma)$ means that $P(Z_n \leqslant x) \to P(Z \leqslant x)$, for every $x \in \mathbb{R}^d$, and a vector $Z \sim N(0, \Sigma)$.

Lemma 3. If $\sqrt{n}\left(\widehat{\theta}_n-\theta_0\right) \rightsquigarrow N_d(0,\Sigma_0)$, for some $\theta_0\in\Theta$ and positive-definite matrix Σ_0, then $\sqrt{n}\left(\eta\left(\widehat{\theta}_n\right)-\eta_0\right) \rightsquigarrow N\left(0,\tau_0^2\right)$, for $\eta_0=\eta(\theta_0)$ and $\tau_0^2=H_\eta(\theta_0,\eta_0)^{-2}H_\theta(\theta_0,\eta_0)^T\Sigma_0 H_\theta(\theta_0,\eta_0)$.

Proof. Abbreviate $\widehat{\eta}=\eta\left(\widehat{\theta}_n\right)$. By the monotonicity of H in its second argument $\Pr(\widehat{\eta}>\eta_0+\epsilon)\leqslant\Pr\left(H\left(\widehat{\theta}_n,\eta_0+\epsilon\right)\leqslant 0\right)$, for any $\epsilon>0$. Since $H\left(\widehat{\theta}_n,\eta_0+\epsilon\right)\to H(\theta_0,\eta_0+\epsilon)$ in probability by continuity of H, and $H(\theta_0,\eta_0+\epsilon)>0$, it follows that $\Pr(\widehat{\eta}>\eta_0+\epsilon)\to 0$. Combined with a similar bound on the left side, this shows that $\widehat{\eta}\to\eta_0$, in probability.

Define a function R by

$$R(\theta,\eta)=\frac{H(\theta,\eta)-H(\theta_0,\eta_0)-H_\theta(\theta_0,\eta_0)^T(\theta-\theta_0)-H_\eta(\theta_0,\eta_0)(\eta-\eta_0)}{\|\theta-\theta_0\|+|\eta-\eta_0|}.$$

Then $R\left(\widehat{\theta},\widehat{\eta}\right)\to 0$ in probability, by the continuous mapping theorem and because H is assumed differentiable. By the definitions of $\widehat{\eta}$ and η_0 we have

$$\begin{aligned}0&=H\left(\widehat{\theta},\widehat{\eta}\right)-H(\theta_0,\eta_0)\\&=H_\theta(\theta_0,\eta_0)^T\left(\widehat{\theta}-\theta_0\right)+H_\eta(\theta_0,\eta_0)(\widehat{\eta}-\eta_0)+R\left(\widehat{\theta},\widehat{\eta}\right)\left(\|\widehat{\theta}-\theta_0\|+|\widehat{\eta}-\eta_0|\right).\end{aligned}$$

It follows that, with $0/0=0$,

$$\sqrt{n}(\widehat{\eta}-\eta_0)\left(H_\eta(\theta_0,\eta_0)+o_P(1)\frac{|\widehat{\eta}-\eta_0|}{\widehat{\eta}-\eta_0}\right)=-H_\theta(\theta_0,\eta_0)^T\sqrt{n}\left(\widehat{\theta}-\theta_0\right)+o_P(1)\sqrt{n}\,\|\widehat{\theta}-\theta_0\|.$$

The result follows by Slutzky's lemma and the fact that $H_\eta(\theta_0,\eta_0)>0$, by assumption.

Next assume that for a given prior distribution on Θ, we obtain a posterior distribution $\Pi(\theta|Y_n)$. Let $\|\cdot\|$ denote the total variation norm. Under the Bernstein-von Mises theorem (cf. e.g., van der Vaart[5], Chapter 10), the posterior distribution permits a normal approximation as assumed in the following lemma.

Lemma 4. If $\|\Pi(\theta|Y_n)-N_d\left(\widehat{\theta}_n,n^{-1}\Sigma_0\right)\|\to 0$ in probability, for estimators $\widehat{\theta}_n$ such that $\sqrt{n}\left(\widehat{\theta}_n-\theta_0\right)\rightsquigarrow N(0,\Sigma_0)$, for some $\theta_0\in\Theta$ and positive-definite matrix Σ_0, then $\|\Pi(\eta(\theta)|Y_n)-N\left(\eta\left(\widehat{\theta}_n\right),n^{-1}\tau_0^2\right)\|\to 0$ in probability, for τ_0 as in the preceding lemma.

Proof. Fix arbitrary $a_2,\ldots,a_d\in\mathbb{R}^d$ such that $H_\theta(\theta_0,\eta_0),a_2,\ldots,a_d$ are orthogonal in $\mathbb{R}^d$ and $a_2,\ldots,a_d$ have norm 1. Define a map $\overline{H}:\Theta\times\mathbb{R}^d\to\mathbb{R}^d$ by

$$\overline{H}(\theta,\eta)=\begin{pmatrix}H(\theta,\eta_1)\\ a_2^T\theta-\eta_2\\ \vdots\\ a_d^T\theta-\eta_d\end{pmatrix}.$$

Then $\overline{H}$ is continuously differentiable with

$$\frac{\partial\overline{H}(\theta,\eta)}{\partial\theta_1\cdots\partial\theta_d}=\begin{pmatrix}H_\theta(\theta,\eta_1)^T\\ a_2^T\\ \vdots\\ a_d^T\end{pmatrix},\qquad \frac{\partial\overline{H}(\theta,\eta)}{\partial\eta_1\cdots\partial\eta_d}=\begin{pmatrix}H_\eta(\theta,\eta_1) & 0 & \cdots & 0\\ 0 & 1 & \cdots & 0\\ \vdots & \vdots & & \vdots\\ 0 & 0 & \cdots & 1\end{pmatrix}.$$

The second matrix at $(\theta_0, \overline{\eta}_0)$, where $\overline{\eta}_0 = \left(\eta(\theta_0), a_2^T\theta_0, ..., a_d^T\theta_0^T\right)$ is the solution to $H(\theta_0, \eta) = 0$, is invertible. By the implicit function theorem there exists an open neighbourhood G of θ_0 and a diffeomorphism $\overline{\eta}: G \to \overline{\eta}(G) \subset \mathbb{R}^d$ such that $\overline{H}(\theta, \overline{\eta}(\theta)) = 0$, for $\theta \in G$. Clearly $\overline{\eta}_1(\theta) = \eta(\theta)$. Abbreviate $\widehat{\overline{\eta}} = \overline{\eta}\left(\widehat{\theta}\right)$ and $\overline{\eta}_0 = \overline{\eta}(\theta_0)$.

By the same method as in the preceding lemma we find that $\sqrt{n}\left(\widehat{\overline{\eta}} - \overline{\eta}_0\right)$ tends in distribution to a normal distribution. For a given bounded set $B \subset \mathbb{R}^d$, the probability of the event $\widehat{\overline{\eta}} + B/\sqrt{n} \subset \overline{\eta}(G)$ therefore tends to one. On this event we have that $\sqrt{n}\left(\overline{\eta}(\theta) - \widehat{\overline{\eta}}\right) \in B$ if and only if $\theta \in \overline{\eta}^{-1}\left(\widehat{\overline{\eta}} + B/\sqrt{n}\right)$, where $\overline{\eta}^{-1}$ is the (ordinary) inverse function of $\overline{\eta}$. By assumption $\theta \,|\, Y_n$ is asymptotically distributed as $\widehat{\theta} + \Sigma_0^{1/2}Z/\sqrt{n}$, for a standard normal vector Z. We thus have that, with ϕ the density of Z,

$$\Pr\left(\sqrt{n}\left(\overline{\eta}(\theta) - \widehat{\overline{\eta}}\right) \in B | Y_n\right) = \Pr\left(\theta \in \overline{\eta}^{-1}\left(\widehat{\overline{\eta}} + B/\sqrt{n}\right) | Y_n\right) = \Pr\left(\widehat{\theta} + \Sigma_0^{1/2}Z/\sqrt{n} \in \overline{\eta}^{-1}\left(\widehat{\overline{\eta}} + B/\sqrt{n}\right)\right) + o_P(1) \tag{A1}$$

$$\begin{aligned} &= \int_{z:\sqrt{n}\left(\overline{\eta}\left(\widehat{\theta}+\Sigma_0^{1/2}z/\sqrt{n}\right) - \widehat{\overline{\eta}}\right) \in B} \phi(z)\,dz + o_P(1) \\ &= \int_B \phi\left(\sqrt{n}\Sigma_0^{-1/2}\left(\overline{\eta}^{-1}\left(\widehat{\overline{\eta}} + y/\sqrt{n}\right) - \widehat{\theta}\right)\right) \left|\frac{dz}{dy}\right| dy + o_P(1), \end{aligned} \tag{A2}$$

where we have made the substitution $\sqrt{n}\left(\overline{\eta}\left(\widehat{\theta} + \Sigma_0^{1/2}z/\sqrt{n}\right) - \widehat{\overline{\eta}}\right) = y$, and $|\,dz/dy\,|$ is (the determinant of) the Jacobian of the transformation. Here

$$\begin{aligned} &\sqrt{n}\Sigma_0^{-1/2}\left(\overline{\eta}^{-1}\left(\widehat{\overline{\eta}} + y/\sqrt{n}\right) - \widehat{\theta}\right) \to \Sigma_0^{-1/2}\left(\overline{\eta}^{-1}\right)'(\eta_0)y = \Sigma_0^{-1/2}\overline{\eta}'(\theta_0)^{-1}y, \\ &\frac{dz}{dy} = \Sigma_0^{-1/2}\left(\overline{\eta}^{-1}\right)'\left(\widehat{\overline{\eta}} + y/\sqrt{n}\right) \to \Sigma_0^{-1/2}\overline{\eta}'(\theta_0)^{-1}. \end{aligned}$$

The convergence is uniform in y ranging over bounded sets. It follows that the integral tends to

$$\int_B \phi\left(\Sigma_0^{-1/2}\overline{\eta}'(\theta_0)^{-1}y\right) \,|\, \Sigma_0^{-1/2}\overline{\eta}'(\theta_0)^{-1} \,|\, dy.$$

We recognise this as the probability of the set B under the normal distribution $N\left(0, \overline{\eta}'(\theta_0)\Sigma_0\overline{\eta}'(\theta_0)^T\right)$. The first marginal of this distribution is the normal distribution in the lemma.

To complete the proof we argue that the convergence is valid in the sense of the total variation distance. The approximation of $\Pi(\theta \in \cdot\,|Y_n)$ with a Gaussian distribution is valid uniformly in B by assumption. Thus the approximation (A1) is valid uniformly, and it suffices to show that (A1) tends to the final equation, uniformly in B. Because the integrand in (A2) converges uniformly for y in bounded sets, the Gaussian probability in (A1) converges uniformly in bounded sets B. Since we can find a sufficiently large compact set that contains most of the mass, this suffices.

As a consequence of the preceding lemma, the posterior median or posterior trimmed mean $\widehat{\eta}$ of $\Pi(\eta(\theta)|Y_n)$ satisfies

$$\sqrt{n}(\widehat{\eta} - \eta(\theta_0)) \rightsquigarrow N\left(0, \tau_0^2\right).$$

Convergence of the posterior mean would require some extra conditions, as the range of $\eta(\theta)$ may be unbounded, depending on H.

Finally given a posterior distribution $\Pi(\theta|Y_n)$, consider the solution $\widetilde{\eta}(Y_n)$ to $\mathrm{E}(H(\theta, \eta)|Y_n) = 0$. Thus we first take the average of the curves $\eta \mapsto H(\theta, \eta)$ relative to the posterior distribution, and then determine a zero. Although in general $\widetilde{\eta}_n$ will be different from any of the preceding, the following lemma shows that it is asymptotically very close to the plug-in estimator using the posterior mean $\mathrm{E}(\theta|Y_n)$ as the estimator of θ.

Strengthen the assumptions on H to: H is bounded and twice continuously differentiable. Write $H_{\theta\theta}$ for the second derivative matrix of H relative to θ.

Lemma 5. Suppose that $n\mathrm{Cov}(\theta|Y_n) \to \Sigma_0$ in probability, and $n^2\mathrm{E}\left(\|\theta-\widehat{\theta}_n\|^4|Y_n\right) = O_P(1)$, for $\widehat{\theta}_n = \mathrm{E}(\theta|Y_n)$. Assume that $\left(\widehat{\theta}_n, \eta\left(\widehat{\theta}_n\right)\right) \to (\theta_0,\eta_0)$ in probability. Then $n\left(\widetilde{\eta}_n - \eta\left(\widehat{\theta}_n\right)\right) \to -\frac{1}{2}H_\eta(\theta_0,\eta_0)^{-1}\mathrm{tr}(H_{\theta\theta}(\theta_0,\eta_0)\Sigma_0)$, in probability.

Proof. Write $\widehat{\eta} = \eta\left(\widehat{\theta}\right)$. By the monotonicity of $\eta \mapsto \mathrm{E}(H(\theta,\eta)|Y_n)$ and the definition of $\widetilde{\eta}$ we have $\Pr(\widetilde{\eta} > \eta_0+\epsilon) \leqslant \Pr(\mathrm{E}(H(\theta,\eta_0+\epsilon)|Y_n) \leqslant 0)$, for every $\epsilon > 0$, which tends to zero since $\mathrm{E}(H(\theta,\eta_0+\epsilon)|Y_n) \to H(\theta_0,\eta_0+\epsilon) > 0$, in probability. Combined with a similar argument for the left side, this shows that $\widetilde{\eta} \to \eta_0$, in probability.

By definition $\mathrm{E}(H(\theta,\widetilde{\eta})|Y_n) = 0$ and $\mathrm{E}\left(\theta-\widehat{\theta}|Y_n\right) = 0$. Then for $\delta > 0$, by the assumed boundedness of H and Markov's inequality,

$$|\,\mathrm{E}\left(H(\theta,\widetilde{\eta})1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right)| = |\,\mathrm{E}\left(H(\theta,\widetilde{\eta})1_{\|\theta-\widehat{\theta}\|>\delta}|Y_n\right)| \leqslant \frac{1}{\delta^4}\mathrm{E}\left(\|\theta-\widehat{\theta}_n\|^4|Y_n\right),$$

$$|\,\mathrm{E}\left(\left(\theta-\widehat{\theta}\right)1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right)| = |\,\mathrm{E}\left(\left(\theta-\widehat{\theta}\right)1_{\|\theta-\widehat{\theta}\|>\delta}|Y_n\right)| \leqslant \frac{1}{\delta^3}\mathrm{E}\left(\|\theta-\widehat{\theta}_n\|^4|Y_n\right).$$

For $\delta = \delta_n \downarrow 0$ slowly enough that $\delta_n^4 n \to \infty$, both expressions are $o_P(n^{-1})$, by the assumption on the fourth moment.

Define $\widehat{R}(\theta,\eta)$ as

$$\widehat{R}(\theta,\eta) = \frac{H(\theta,\eta) - H\left(\widehat{\theta},\widehat{\eta}\right) - P_2(\theta,\eta)}{\|\theta-\widehat{\theta}\|^2 + |\eta-\widehat{\eta}|^2},$$

where $P_2(\theta,\eta)$ is the second order Taylor polynomial of H around $\left(\widehat{\theta},\widehat{\eta}\right)$, i.e.

$$P_2(\theta,\eta) = H_\theta\left(\widehat{\theta},\widehat{\eta}\right)^T\left(\theta-\widehat{\theta}\right) + H_\eta\left(\widehat{\theta},\widehat{\eta}\right)(\eta-\widehat{\eta})$$
$$+\frac{1}{2}\left(\theta-\widehat{\theta}\right)^T H_{\theta\theta}\left(\widehat{\theta},\widehat{\eta}\right)\left(\theta-\widehat{\theta}\right) + \left(\theta-\widehat{\theta}\right)^T H_{\theta\eta}\left(\widehat{\theta},\widehat{\eta}\right)(\eta-\widehat{\eta}) + \frac{1}{2}H_{\eta\eta}\left(\widehat{\theta},\widehat{\eta}\right)(\eta-\widehat{\eta})^2.$$

Since $H\left(\widehat{\theta},\widehat{\eta}\right) = 0$ by definition of $\widehat{\eta}$,

$$\mathrm{E}\left(H(\theta,\widetilde{\eta})1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right) = \mathrm{E}\left(P_2(\theta,\widetilde{\eta})1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right) + \mathrm{E}\left(R(\theta,\widetilde{\eta})\left(\|\theta-\widehat{\theta}\|^2 + |\widetilde{\eta}-\widehat{\eta}|^2\right)1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right).$$

By the preceding the left side is $o_P(n^{-1})$, if $\delta \downarrow 0$ sufficiently slowly. Since the supremum of R over a shrinking neighbourhood of (θ_0,η_0) tends to zero, we can bound R out of the conditional expectation and that the remainder term gives a contribution of the order $o_P(n^{-1})$. The term involving P_2 contributes five terms. The term that is linear in $\theta-\widehat{\theta}$ was already seen to be $o_P(n^{-1})$, while the mixed term in $\theta-\widehat{\theta}$ and $\widetilde{\eta}-\widehat{\eta}$ is even smaller due to the presence of $\widetilde{\eta}-\widehat{\eta}$. Th quadratic term in $\widetilde{\eta}-\widehat{\eta}$ is negligible relative to the linear term in this variable. Thus rearranging the terms yields

$$(\widetilde{\eta}-\widehat{\eta})\left(H_\eta(\theta_0,\eta_0) + o_P(1)\right) = -\frac{1}{2}\mathrm{E}\left(\left(\theta-\widehat{\theta}\right)^T H_{\theta\theta}\left(\widehat{\theta},\widehat{\eta}\right)\left(\theta-\widehat{\theta}\right)1_{\|\theta-\widehat{\theta}\|\leqslant\delta}|Y_n\right) + o_P\left(n^{-1}\right).$$

The indicator within the conditional expectation can be removed, and $H_{\theta\theta}\left(\widehat{\theta},\widehat{\eta}\right)$ replaced by $H_{\theta\theta}(\theta_0,\eta_0)$, at the cost of another $o_P(n^{-1})$-term, The expectation can next be computed as the trace of the matrix $H_{\theta\theta}(\theta_0,\eta_0)\mathrm{Cov}(\theta|Y_n)$. This matrix times n tends to $H_{\theta\theta}(\theta_0,\eta_0)\Sigma_0$, by assumption.

The intended application is to find a "control limit" η such that $\Pr_\theta(Y_T \geqslant \gamma | Y_0 \geqslant \eta) = q$. If θ is unknown, it is replaced by a data-based approximation. The following corollary shows that estimation leads to a random error of order $O_P\left(n^{-1/2}\right)$ in the control level q.

Corollary 1. *Assume that the conditions of the preceding lemmas hold and let $\widehat{\eta}$ be either the plug-in estimator $\eta\left(\widehat{\theta}\right)$, or the median or trimmed mean of the posterior distribution of $\eta(\theta)$, or the solution to* $\mathrm{E}(H(\theta,\eta)|Y_n) = 0$. *Then* $\sqrt{n}(H(\theta_0,\widehat{\eta}) - H(\theta_0,\eta_0)) \rightsquigarrow N\left(0, H_\theta(\theta_0,\eta_0)^T \Sigma_0 H_\theta(\theta_0,\eta_0)\right)$.

Proof. This is an immediate consequence of the three lemmas and the Delta-method.