# Monitoring of Safety and Effectiveness of COVID-19 Therapeutics

## Sentinel Sequential Surveillance Master Protocol

John Connolly, ScD,[1] Marie Bradley, PhD, MPharm, MScPH,[2] José J. Hernández-Muñoz, RPh, MPH, MSc, PhD,[2] Sruthi Adimadhyam, PhD, MS,[1] James Antoon, MD, PhD, MPH,[3] Maya Beganovic, PharmD, MPH, BCIDP,[2] Patricia Bright, PhD, MSPH,[2] Sarah Dutcher, PhD, MS,[2] Meredith Epperson, MScPH,[1] Efe Eworuke, PhD,[2] Joy Kolonoski, MPH,[1] Yong Ma, PhD, MS,[2] Andrew Petrone, MPH,[1] Judith C. Maro, PhD, MS,[1] Sukhminder Sandhu, PhD, MPH, MS[2]

Affiliations: 1. Department of Population Medicine, Harvard Pilgrim Health Care Institute and Harvard Medical School, Boston, MA; 2. Office of Surveillance and Epidemiology, Center for Drug Evaluation and Research, US Food and Drug Administration, Silver Spring, MD; 3. Department of Pediatrics, Monroe Carell Jr Children's Hospital at Vanderbilt and Vanderbilt University Medical Center, Nashville, TN

Version 1.1

January 10, 2024

**The Sentinel System is sponsored by the U.S. Food and Drug Administration (FDA) to proactively monitor the safety of FDA-regulated medical products and complements other existing FDA safety surveillance capabilities. The Sentinel System is one piece of FDA's Sentinel Initiative, a long-term, multi-faceted effort to develop a national electronic system. Sentinel Collaborators include Data and Academic Partners that provide access to healthcare data and ongoing scientific, technical, methodological, and organizational expertise. The Sentinel Coordinating Center is funded by the FDA through the Department of Health and Human Services (HHS) Contract number 75F40119D10037.**

# Disclaimer

## For Patients and Consumers

- The information contained on this website is provided as part of FDA's commitment to place knowledge acquired from Sentinel in the public domain as soon as possible.
- Much of the content on this site is technical and intended for use by scientists in various areas of expertise.
- **The fact that FDA requests and receives data on a particular product through Sentinel does not necessarily mean there is a safety issue with the product.**
- FDA may access the data available through Sentinel for a variety of reasons beyond assessing potential safety risks for a specific product. Some examples include determining a rate or count of an identified health outcome of interest, examining medical product use, or seeking to better understand the capabilities of the Sentinel System.
- When evaluating a potential safety issue, FDA scientists consider the data obtained through Sentinel with information from various other data sources, such as adverse event reports, published study results, and clinical trials, to help make the most informed decisions possible.
- **FDA communicates its interpretation of Sentinel activities through existing channels, such as FDA's press announcements, MedWatch Alerts, and Drug Safety Communications, rather than on this website.**
- Information from this site should not affect your use of a medical product in any way. Patients who have questions about the use of a medical product should contact their health care professional.

## For Industry

The information contained on this website is provided as part of FDA's commitment to place knowledge acquired from Sentinel in the public domain as soon as possible. To most effectively interpret results from observational studies, it is important to consider not only the studies that supported a hypothesis, but also the studies that did not. The website serves as a public data repository that archives all the activities of Sentinel and provides important context to those seeking to understand the significance of any specific activity. This information is being provided to the public in the interest of transparency and for purposes of demonstrating the extent of use and the various ways FDA is utilizing the Sentinel System. While the data posted here may contribute to important overall conclusions, FDA relies on other mechanisms for communicating such conclusions to the public.

When reviewing this information please be aware that there are times when FDA may access the data available through Sentinel for a variety of reasons beyond seeking direct access to information that can help assess potential safety risks for a specific product. Some examples include determining a rate or count of an identified health outcome of interest, examining medical product use, exploring the feasibility of future, more detailed analyses within Sentinel, and seeking to better understand the capabilities of Sentinel.

# Sequential Surveillance Master Protocol

## 1 Table of Contents

## History of Modifications

| Version | Date | Modification | Author |
|---|---|---|---|
| 1.0 | 11/17/2023 | Original Version | Sentinel Operations Center |
| 1.1 | 1/10/2024 | First Revision | Sentinel Operations Center |

# 2 Introduction and Objectives

The goal of post-market sequential surveillance is to prospectively identify statistically significant safety signals as quickly as possible by repeatedly analyzing data as they accrue.[1] The Sentinel Distributed Database is an especially valuable resource for sequential surveillance because it can be periodically leveraged to create large cohorts of insured patients with curated data on inpatient and outpatient visits as well as pharmacy dispensings.[2] Prospective sequential analyses have previously been conducted in the Sentinel Distributed Database.[3–6] We expect prospective surveillance to be increasingly necessary given the need to rapidly conduct safety assessments as with the recent COVID-19 pandemic and approval of novel treatments. Therefore, our objective in this master protocol is to provide guidance for future sequential analyses in the Sentinel Distributed Database.

# 3 Methods

## 3.1 Exposures

This protocol is intended to be generally applicable to any medical product(s) of interest to investigators; however, this protocol does not provide guidance specific to the prospective surveillance of vaccines. For further information on sequential analysis of vaccines, see here. Exposures suitable for sequential analysis should be well-measured in the database and used frequently.

Well-measured exposures in the Sentinel Distributed Database include drugs dispensed through outpatient pharmacies or administered through a medical procedure. Frequent utilization of the exposure is necessary to accrue enough outcomes to complete the sequential analysis with reasonable statistical power in a reasonable amount of time. This is a point of special consideration when studying newly marketed drugs which do not have rapid uptake. To quantify whether this condition is met, investigators should conduct pre-analysis calculations which consider the effect size of interest and desired statistical power to determine the length of the sequential analysis. Further detail on these calculations is provided below in the Sequential Surveillance Parameters section.

## 3.2 Outcomes

In contrast to outcome-agnostic methods like TreeScan,[7] sequential surveillance requires pre-specification of an outcome of interest. The outcome may be selected based on existing evidence for an exposure-outcome association, such as adverse event reports in the Food and Drug Administration's Adverse Event Reporting System (FAERS). Alternatively, an outcome may be chosen based on clinical or theoretical grounds alone. As such, sequential surveillance is a flexible methodology that can be applied within analyses which are designed to more rigorously evaluate known potential safety signals or within analysis designed to generate new safety signals.[8]

As with exposures, outcomes which are most suitable for a sequential analysis are those that are both well-measured and sufficiently common. Well-measured outcomes in claims data are those for which validation studies are available. The care setting in which outcomes are identified influences how well they are measured; for example, outcomes which require inpatient care are expected to be more reliably measured than those that do not. Whether an outcome is sufficiently common to complete the sequential analysis in a timely fashion can be quantified through pre-analysis calculations, which are discussed in further detail in the Sequential Surveillance Parameters section.

## 3.3 Study Design

Sequential statistical analysis is compatible with multiple study designs. The appropriate study design depends primarily on the exposures and outcomes of interest. The Sentinel routine analytic tools currently support two study designs which can be used to perform prospective sequential surveillance.

### 3.3.1 Active Comparator New User Design

The active comparator new user cohort design attempts to mimic an analogous randomized controlled trial and has several advantages for prospective surveillance. The selection of an appropriate active comparator for the exposure of interest reduces both measured and unmeasured confounding by indication.[9] When identifying or evaluating safety signals, confounding control is important to reduce false positives. Restriction to new users of both exposures of interest additionally provides a clearly defined "time zero" which is important for avoiding time-related biases related to covariate measurement as well as proper allocation of follow-up time.[10] The active comparator new user design is also flexible, in that it can be used to study exposures which are administered for either a short or long duration and outcomes that have both short and long induction periods.

An ideal active comparator is one which is used by a similar patient population as the exposure of interest. For example, a treatment approved for a similar indication, in the same therapeutic class, or used at the same point in the course of disease. Investigators can quantitatively evaluate a potential active comparator based on measured covariate balance in the population of interest prior to beginning sequential surveillance. When using an active comparator, investigators should carefully consider whether the contrast implied by the selected comparison reflects the scientific question of interest.

To maximize confounding control in a sequential analysis using an active comparator new user design, investigators should consider applying propensity score (PS) methods. In the Sentinel context, statistical methods for sequential analysis are compatible with PS matching and PS stratification. High dimensional PS (HDPS) adjustment may be warranted in settings where confounding is of particular concern.[11] Sequential statistics can be applied to either PS matched or PS stratified analyses. For further discussion on how to incorporate the chosen PS adjustment method, see the Sequential Surveillance Parameters section below.

Between-person designs, like the active comparator new user approach described above, are not without limitations. When conducting prospective safety monitoring of newly marketed medical products, an ideal active comparator may not be available. For example, potential comparators may be infrequently utilized, indicated for use at different stages of disease severity, or the exposure of interest may be the first in its class. It may also be challenging to find concurrent comparators when studying newly marketed medical products which rapidly replace the existing standard of care; this is of importance when there are concurrent secular trends in the outcome which could lead to confounding by calendar time. If a suitable active comparator does exist, interpretation may be challenging if the comparator also has a relationship to the outcome. For example, if both exposure and comparator raise the risk of the outcome of interest, but the comparator raises it to a greater degree, that could obscure a potential safety signal. When there is no single ideal active comparator, multiple comparator groups may be considered within the same analysis. If an active comparator new user design is not feasible, investigators should consider a self-controlled design.

### 3.3.2 Self-Controlled Risk Interval Design

The self-controlled design supported by the Sentinel routine analytic tools is known as the self-controlled risk interval (SCRI) design. In the SCRI design, each patient serves as their own

comparator. Therefore, the primary benefit of the SCRI design is that it controls for all time-invariant confounders through self-matching.[12] This is especially advantageous when conducting sequential surveillance of exposures that have no suitable active comparator or outcomes which are expected to be subject to strong confounding by indication. The SCRI design is most suitable for studying short-term exposures and outcomes with a well-defined date of onset and which occur relatively soon after exposure. Therefore, the SCRI design has previously been used in Sentinel to conduct sequential surveillance of vaccines for outcomes such as venous thromboembolism (VTE) and intussusception.[13,14]

Selection of an appropriate risk and control window is critical. The risk window is typically specified as the period immediately following exposure initiation and should last as long as the exposure can plausibly increase outcome risk. The control window may be specified during any time outside the risk window; however, it is typically chosen to follow the risk window to avoid issues with what is known as the "healthy vaccinee" effect.[15] This bias may occur when pre-exposure time is used for the control window and the occurrence of the outcome influences the probability of exposure. Additionally, the risk and control windows may differ in length; however, the specified lengths must not vary across patients.

Although there are inherent strengths of the SCRI design, there are additional limitations that should be considered when choosing a design for the sequential surveillance. First, SCRI is less flexible than the new user active comparator design and is not appropriate for long-duration exposures or outcomes which have either an insidious onset or long induction time. In addition, although the self-controlled design addresses time invariant confounders, any time varying confounding between the risk and control windows must be explicitly modeled when using the SCRI design. Time varying confounding may be particularly problematic when studying treatments given to patients who are acutely ill.

## 3.4 Sequential Surveillance Parameters

In the post-marketing safety context, the goal of sequential analyses is to identify statistically significant increases in the risk of an adverse event as quickly as possible. Recent methodological advances allow for sequential analyses using exact statistical methods which can be performed in either a continuous or group fashion.[1] In a continuous sequential analysis, a new statistical test is performed each time new data accrue. In contrast, a group sequential analysis involves a new statistical test only after an investigator-specified amount of data have accrued. Whether a continuous or group sequential design is preferable depends on the objective of the analysis. Group sequential analysis is preferable when the objective is to minimize expected sample size, a common concern in randomized trials where gathering additional data may come with significant costs. When the objective is to minimize the time needed to reject the null hypothesis of no increased risk, continuous testing is always preferable to a group testing.[16] Therefore, this protocol focuses exclusively on the continuous version of sequential analyses.

### 3.4.1 Pre-analysis Parameter Specification

Prior to performing a sequential analysis, investigators must specify several parameters. Sequential statistics are flexible and can be applied to a variety of analyses with competing goals. When applied specifically to post-market safety monitoring within the Sentinel Distributed Database, many parameters have a clearly preferred specification. For example, investigators must decide whether their data arise from a binomial or Poisson distribution. This master protocol focuses exclusively on the binomial version of sequential statistics as they are more compatible with Sentinel's distributed analytic tools used to implement the two study designs discussed above.[3] Next, investigators must specify their null and alternative hypotheses. As the focus of this master protocol is to describe analyses identifying statistically significant increases

in risk, we assume a null hypothesis $H_0$: Risk Ratio ≤ 1 and an alternative hypothesis $H_1$: Risk Ratio > 1. Along similar lines, investigators must decide if they prefer to minimize the expected time signal or the expected sample size. In this protocol, the focus is to minimize the expected time to signal because, as mentioned above, for sequential analyses of post-market safety using observational data investigators should always prefer to minimize the expected time to signal.

Investigators must also specify an alpha level and alpha spending plan. The specified alpha level represents the Type I error probability. By convention this is typically specified as 0.05; however, investigators should consider whether a different alpha level is appropriate for their analysis based on the context-specific costs of Type I vs. Type II errors. The alpha spending plan outlines how much of the allotted alpha level is used at each sequential test, depending on the goals of the analysis. A recently developed alpha spending method can be selected to optimize (i.e., minimize) the expected time to detect a safety signal.[17] This approach has been shown to provide shorter expected time to signaling than alternative alpha spending approaches, and it adaptively spends the appropriate amount of alpha at each test. Therefore, this protocol will focus exclusively on the optimal alpha spending plan.

### 3.4.2 Pre-analysis Calculations

After specifying the above parameters, investigators should implement the pre-analysis calculation mentioned in the Exposures and Outcomes sections above. Specifically, investigators should calculate the shape of the alpha spending function, the expected time to signal, and the maximum length of the analysis. The expected time to signal is defined as the expected number of observed outcomes required to reject the null hypothesis. The maximum length of the analysis is defined as the number of observed outcomes required to completely spend all alpha when the null hypothesis is not rejected.

This calculation can be performed using the `Optimize.Binomial` function in the R package 'Sequential'. To perform this calculation, investigators must provide a quantity to optimize (i.e., expected time to signal), an alpha level, a risk ratio of interest, the desired statistical power, whether the test is one- or two-tailed, and the probability that an outcome occurs in an exposed patient under the null hypothesis. The first two parameters have been discussed above. The target risk ratio and desired statistical power should be chosen based on a balance between clinically meaningful changes in the relative risk of the outcome and practical considerations regarding the required sample size, as discussed further below. Whether the test is one- or two-tailed depends on the null and alternative hypotheses. For the assumed null and alternative hypotheses discussed in this protocol ($H_0$: Risk Ratio ≤ 1, $H_1$: Risk Ratio > 1), the one-tailed option is preferred. The expected probability of an outcome arising in an exposed versus comparator patient under the null hypothesis (or, analogously in the SCRI design, from within the risk versus the control window) depends on the study design. For example, if an active comparator new-user design used 1:1 PS matching, then the expected probability that an observed outcome occurs in the exposed group when the true risk ratio is one is 0.5. We would expect the same probability if, in an SCRI design, the risk and control windows are of equal length. If a different PS matching ratio was used, or the risk and control windows were not of equal length, this probability should be appropriately adjusted.

Table 1 displays the expected time to signal and maximum length of the analysis (both measured in the total number of observed outcomes) for selected combinations of a target risk ratio and desired statistical power with the following parameter specifications: optimizing for a minimum expected time to signal, the probability of an outcome coming from the exposed group of 0.5, an alpha level of 0.05, and a one-tailed hypothesis test. As an example, Figure 1 displays the shape of the optimized alpha spending function when the target risk ratio is 3 and the target power is 0.9.

**Table 1: Expected time to signal and maximum analysis length for various target risk ratios and power**

| Target Risk Ratio | Target Power | Expected time to signal[1] | Maximum length of analysis[2] |
|---|---|---|---|
| 1.5 | 0.8 | 94.0 | 177 |
| 2 | 0.8 | 33.1 | 77 |
| 3 | 0.8 | 15.2 | 32 |
| 1.5 | 0.9 | 119.8 | 232 |
| 2 | 0.9 | 40.5 | 97 |
| 3 | 0.9 | 18.2 | 43 |

[1]Expected time to signal is the expected number of observed outcomes required to reject null hypothesis
[2]Maximum length of the analysis is the number of observed outcomes required to spend all alpha when null hypothesis is not rejected

**Figure 1: Alpha spending function for one-tailed test with an alpha = 0.05, target risk ratio = 3, target power = 0.9**

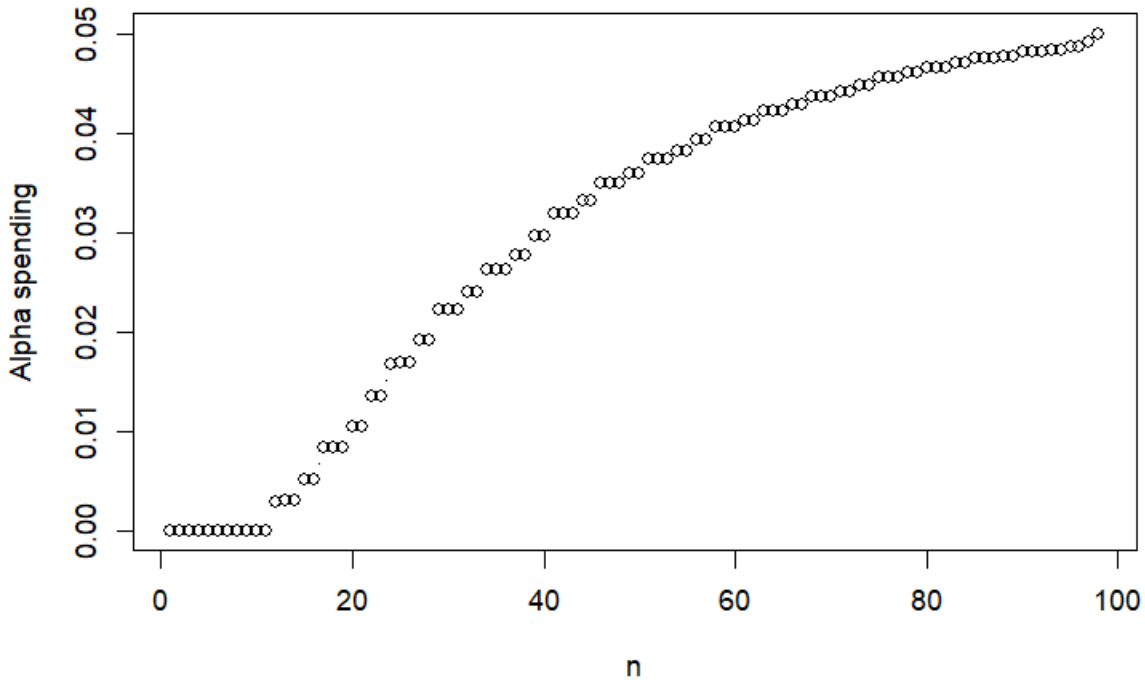


Table 1 shows how the expected time to signal and maximum length of the analysis increase as the desired power increases and the target risk ratio decreases. For the combination of parameters used to generate Figure 1, no alpha is spent until the 10th outcome. This implies that, for this specific analysis, the null hypothesis cannot be rejected until at least the 10th outcome has occurred. Performing the pre-analysis derivation of the optimal alpha spending plan may be computationally intensive when the maximum length of the analysis is above 200 outcomes.[17] Fortunately, this is not a concern for most practical combinations of a target risk ratio and desired statistical power.

Along similar lines, investigators should carefully consider whether the maximum length of the analysis, which is measured in the number of outcomes, can be observed in a reasonable length

of calendar time before beginning the sequential analysis. If studying a rare outcome, or a rare exposure, it may take an unreasonably long amount of calendar time to accumulate enough outcomes to complete the sequential analysis. In such cases, investigators may consider increasing the target risk ratio or decreasing the target power. They may also consider using a different outcome definition (for example, including additional care settings) or a different outcome altogether.

### 3.4.3 Sequential Testing

In a continuous sequential analysis where the goal is to minimize the expected time to signal, a new test or "look" should be performed each time new data (i.e., outcomes) accrue. In the Sentinel Distributed Database, new data accrue at each Data Partner data refresh. Each sequential test can be performed using the R package 'Sequential'. The cumulative number of exposed and comparator cases (analogously, cases from within the risk and control windows for an SCRI design) at the time of each look is used to calculate a log-likelihood ratio test statistic. This test statistic is then compared to a threshold, or critical value derived from the optimized alpha spending function. If the observed log-likelihood test statistic exceeds the critical value, the null hypothesis is rejected at that look. If the null hypothesis is never rejected, the analysis continues until all alpha is spent.

While the parameter values required to perform the analysis are specified before the analysis begins, they may be adjusted at each look if necessary. For example, the probability of an exposed outcome under the null may be changed at each look, which may be necessary for variable ratio PS matched studies. Investigators may also deviate from the original alpha spending plan and spend a custom amount at a given look. This option allows investigators to complete the sequential analysis earlier than planned while maintaining the desired Type I error probability. This option may be especially useful when there are fewer observed outcomes than expected, perhaps due to slow uptake of the exposure of interest. While the formal sequential analysis ends upon rejection of the null or spending all alpha, investigators may elect to continue safety monitoring outside the formal sequential analysis.

## 3.5 Data Quality

Sequential analyses assume data used in a prior test do not change upon subsequent tests. When conducting a sequential analysis in claims data, this assumption is unlikely to completely hold.[18] A small proportion of claims may appear, disappear, or change dates after a Data Partner refresh due to the inherent lag time between the occurrence of a healthcare claim and its availability for research in the Sentinel Distributed Database. Such changes can impact any aspect of the analysis, including exposure, covariates, outcome, and cohort eligibility. While unsettled claims are not unique to sequential analyses, the issue must be addressed because sequential analyses span across multiple Data Partner refreshes. Of particular concern may be claims from inpatient care settings, which take longer to settle than other types and are often used to define outcomes.[19]

The probability of including unsettled claims can be reduced through careful specification of the study period. The Sentinel tools allow investigators to select the last day of follow-up based on the relative completeness of data at that site. Specifically, the tools define the final day of follow-up as the end of the latest month where the record count is greater than or equal to 80% of the count observed in the prior month across all encounter types (e.g., ambulatory, inpatient, emergency department, etc.). Table 2 provides an example of the relative completeness algorithm.

**Table 2: Relative completeness algorithm to determine Data Partner End Date**

| Month | Number of Encounters in Month | Number of Encounters in Month / Number of Encounters in Prior Month |
|---|---|---|
| January 2022 | 10,000 | N/A |
| February 2022 | 12,000 | 120% |
| March 2022 | 11,500 | 96% |
| April 2022 | 9,200 | 80% |
| May 2022 | 500 | 6% |

According to the 80% completeness algorithm, the final day of follow-up for the hypothetical Data Partner displayed in Table 2 algorithm would be April 30th, 2022, and all episodes would be censored on that calendar date. If investigators feel this algorithmically defined date is not sufficient to allow claims to settle, the tools allow them to determine the end of follow-up for each look at each Data Partner as a user-defined calendar date. For example, a more conservative threshold in the above example would be March 31st, 2022, which allows more time for claims to settle but induces a lag period which may increase the calendar time to signal.

Ultimately, unsettled claims are an inherent limitation of claims data sources which cannot be completely avoided regardless of the method used to the determine the end of follow-up. Therefore, we recommend fixing data used in prior sequential tests and not allowing any changes as more data accumulate. This approach ensures a key assumption of sequential statistics is met. To assess the impact of fixing data on the number of outcomes, investigators may consider running a separate analysis utilizing data from the entire study period in a single look and comparing the outcome count with that from the multi-look sequential analysis.

## 3.6 COVID-19 Case Study

For the first application of this master protocol, we prefer an exposure used in the treatment of COVID-19. Therefore, we propose a cohort study comparing adverse events in users of corticosteroids and azithromycin versus azithromycin alone in patients recently diagnosed with COVID-19 in an outpatient care setting.

For the outcome of interest, we prefer a previously identified adverse event associated with corticosteroid use but not azithromycin. We also prefer an adverse event which is reliably identified in administrative claims data. Therefore, we propose as the outcome of interest severe hyperglycemia, defined as a diagnosis of hyperglycemia in the inpatient or emergency department care setting. Corticosteroids are known to impair glucose control and exacerbate hyperglycemia in patients with pre-existing diabetes.[20] Additionally, hyperglycemia and associated complications (e.g., decreased carbohydrate and glucose tolerance, diabetes mellitus, glycosuria, etc.) are labeled adverse events for corticosteroids.[21–23] Given the choice of outcome,, we will limit our study population to adults with Type 2 diabetes. We will assume all outcomes arise from a binomial distribution and specify a one-sided null hypothesis $H_0$: Risk Ratio $\leq 1$ and an alternative hypothesis $H_1$: Risk Ratio $> 1$. Our goal will be to minimize the expected time to signal, and we will use an optimal alpha spending plan.

Further details will be provided during the implementation of the case study. We acknowledge that some aspects of the proposed case study may change in response to developments in the COVID-19 pandemic, changing treatment practices, and pre-analysis feasibility assessments.

# 4 References


1. R Silva I, Maro J, Kulldorff M. Exact sequential test for clinical trials and post-market drug and vaccine safety surveillance with Poisson and binary data. *Stat Med* 2021; **40**: 4890–4913. doi:10.1002/sim.9094.

2. Platt R, Brown JS, Robb M, *et al.* The FDA Sentinel Initiative - An Evolving National Resource. *N Engl J Med* 2018; **379**: 2091–2093. doi:10.1056/NEJMp1809643.

3. Fireman B, Toh S, Butler MG, *et al.* A protocol for active surveillance of acute myocardial infarction in association with the use of a new antidiabetic pharmaceutical agent. *Pharmacoepidemiol Drug Saf* 2012; **21 Suppl 1**: 282–290. doi:10.1002/pds.2337.

4. Chrischilles EA, Gagne JJ, Fireman B, *et al.* Prospective surveillance pilot of rivaroxaban safety within the US Food and Drug Administration Sentinel System. *Pharmacoepidemiol Drug Saf* 2018; **27**: 263–271. doi:10.1002/pds.4375.

5. Leonard C, Hennessy, S et al. Prospective Routine Observational Monitoring of Mirabegron. : 18.

6. Toh S, Reichman ME, Graham DJ, *et al.* Prospective Postmarketing Surveillance of Acute Myocardial Infarction in New Users of Saxagliptin: A Population-Based Study. *Diabetes Care* 2018; **41**: 39–48. doi:10.2337/dc17-0476.

7. Kulldorff M, Dashevsky I, Avery TR, *et al.* Drug safety data mining with a tree-based scan statistic. *Pharmacoepidemiol Drug Saf* 2013; **22**: 517–523. doi:10.1002/pds.3423.

8. Martin D, Gagne JJ, Gruber S, *et al.* Sequential surveillance for drug safety in a regulatory environment. *Pharmacoepidemiol Drug Saf* 2018; **27**: 707–712. doi:10.1002/pds.4407.

9. Lund JL, Richardson DB, Stürmer T. The active comparator, new user study design in pharmacoepidemiology: historical foundations and contemporary application. *Curr Epidemiol Rep* 2015; **2**: 221–228. doi:10.1007/s40471-015-0053-5.

10. Hernán MA, Robins JM. Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available. *Am J Epidemiol* 2016; **183**: 758–764. doi:10.1093/aje/kwv254.

11. Rassen JA, Schneeweiss S. Using high-dimensional propensity scores to automate confounding control in a distributed medical product safety surveillance system. *Pharmacoepidemiol Drug Saf* 2012; **21 Suppl 1**: 41–49. doi:10.1002/pds.2328.

12. Baker MA, Lieu TA, Li L, *et al.* A vaccine study design selection framework for the postlicensure rapid immunization safety monitoring program. *Am J Epidemiol* 2015; **181**: 608–618. doi:10.1093/aje/kwu322.

13. Yih WK, Greene SK, Zichittella L, *et al.* Evaluation of the risk of venous thromboembolism after quadrivalent human papillomavirus vaccination among US females. *Vaccine* 2016; **34**: 172–178. doi:10.1016/j.vaccine.2015.09.087.

14. Yih WK, Lieu TA, Kulldorff M, *et al.* Intussusception risk after rotavirus vaccination in U.S. infants. *N Engl J Med* 2014; **370**: 503–512. doi:10.1056/NEJMoa1303164.

15. Hawken S, Potter BK, Little J, *et al.* The use of relative incidence ratios in self-controlled case series studies: an overview. *BMC Med Res Methodol* 2016; **16**: 126. doi:10.1186/s12874-016-0225-0.

16. Silva IR, Kulldorff M. Continuous versus group sequential analysis for post-market drug and vaccine safety surveillance. *Biometrics* 2015; **71**: 851–858. doi:10.1111/biom.12324.

17. Silva IR, Kulldorff M, Katherine Yih W. Optimal alpha spending for sequential analysis with binomial data. *Journal of the Royal Statistical Society Series B, Statistical methodology* 2020; **82**: 1141–1164. doi:10.1111/rssb.12379.

18. Maro JC, Eworuke E, Hou L, *et al.* Conducting prospective sequential surveillance in real-world dynamic distributed databases. *Pharmacoepidemiol Drug Saf* 2020; **29**: 1331–1335. doi:10.1002/pds.5002.

19. Yih WK, Kulldorff M, Sandhu SK, *et al.* Prospective influenza vaccine safety surveillance using fresh data in the Sentinel System. *Pharmacoepidemiol Drug Saf* 2016; **25**: 481–492. doi:10.1002/pds.3908.

20. Li J-X, Cummins CL. Fresh insights into glucocorticoid-induced diabetes mellitus and new therapeutic directions. *Nat Rev Endocrinol* 2022; **18**: 540–557. doi:10.1038/s41574-022-00683-6.

21. DailyMed - DEXAMETHASONE tablet. Available at: https://dailymed.nlm.nih.gov/dailymed/lookup.cfm?setid=bf8acfdb-8dc8-4a69-90b0-33c680c421d6. Accessed August 31, 2023.

22. DailyMed - PREDNISONE tablet. Available at: https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=10fe5a3b-84dc-4600-87c2-b80c97ce18cf. Accessed August 31, 2023.

23. DailyMed - HYDROCORTISONE tablet. Available at: https://dailymed.nlm.nih.gov/dailymed/lookup.cfm?setid=b8437441-a4a6-4fb5-82eb-e49af69704f4. Accessed August 31, 2023.