The Complementary Components of the U.S. Vaccine Safety Monitoring System

Jeffrey S. Morris, Ph.D.
Perelman School of Medicine
APPC Distinguished Research Fellow
University of Pennsylvania

Preface by Kathleen Hall Jamieson, Ph.D.
Director, Annenberg Public Policy Center

A report of the Annenberg Public Policy Center

December 2025

No. 3 in the **APPC Vaccination Communication and Fact-Checking Toolkit** series

Annenberg
PUBLIC POLICY CENTER
UNIVERSITY of PENNSYLVANIA

The Complementary Components of the U.S. Vaccine Safety Monitoring System

Analytical Considerations, Limitations, and Potential Improvements

A report of the Annenberg Public Policy Center

December 2025

Jeffrey S. Morris, Ph.D.
George S. Pepper Professor of Public Health and Preventive Medicine
Director, Division of Biostatistics
Department of Biostatistics, Epidemiology and Informatics
Perelman School of Medicine
University of Pennsylvania
Philadelphia, PA

THE ANNENBERG PUBLIC POLICY CENTER
Kathleen Hall Jamieson, Ph.D.
Patrick E. Jamieson, Ph.D.
Shawn Patterson Jr., Ph.D.

This is the third in a series of white papers that comprise the APPC Vaccination Communication and Fact-Checking Toolkit.

Table of Contents

Acknowledgments

This report includes the input and analysis of the following individuals. Find related reports online at www.annenbergpublicpolicycenter.org.

Editorial and Graphic Design Team

Sofie Adams
Research Administrative Coordinator, APPC

Zachary Reese
Senior Designer, APPC

Karen Riley
Director of Events and Facilities, APPC

Michael Rozansky
Director of Communications, APPC

Lana Xu
Administrative Coordinator, APPC

APPC Vaccination Communication and Fact-Checking Toolkit Series

1) Minimizing public susceptibility to misconceptions about the effects of vaccination: Vaccine Adverse Event Reporting System (VAERS) (May 2023)

2) Reducing Susceptibilities to Misconceptions About Vaccination During Pregnancy: RSV (August 2023)

3) The Complementary Components of the U.S. Vaccine Safety Monitoring System (December 2025)

RECOMMENDED CITATION

The Annenberg Public Policy Center. 2025. "The Complementary Components of the U.S. Vaccine Safety Monitoring System."

Preface

How the Food and Drug Administration (FDA) and the Centers for Disease Control and Prevention (CDC) determine whether a vaccine has resulted in the death of a vaccine recipient became the focus of renewed scrutiny by the press and public last month. In November 2025, one of the nation's chief vaccine regulators, the FDA's Chief Medical and Scientific Officer Vinay Prasad, M.D., M.P.H., reported in a memo that although the FDA had "never publicly admitted it," children in the United States had died as a result of COVID-19 vaccination (Prasad 2025).

The public health community, press, and public will know more about how FDA researchers arrived at this conclusion when the sources and methods they used to tie COVID-19 vaccination to the deaths are revealed. What is known at this point is that the part of the vaccine safety system mentioned in the memo – the Vaccine Adverse Event Reporting System, or VAERS – is a repository for unconfirmed reports that relies on voluntary submissions from individuals, families, healthcare professionals and others, and serves as an early-warning system.

Other resources within the vaccination monitoring system increase researchers' ability to determine whether the signal of a possible concern registered in VAERS or in one of the other sentinel entities reflects an actual concern. Such documentation is a multifaceted process involving data not found in VAERS. Paul Offit, M.D., an infectious disease specialist at Children's Hospital of Philadelphia, indicates that linking vaccination to a death requires an abundance of evidence including autopsy reports and medical records that exclude other causes and show whether the deceased had COVID-19 (Szabo 2025).

Because the Prasad statement focused national attention on the important question of how the federal government assesses the safety of vaccines, this is an opportune time to heighten understanding of the range of safety monitoring entities at the ready in the federal government and the complementarity of central ones among them.

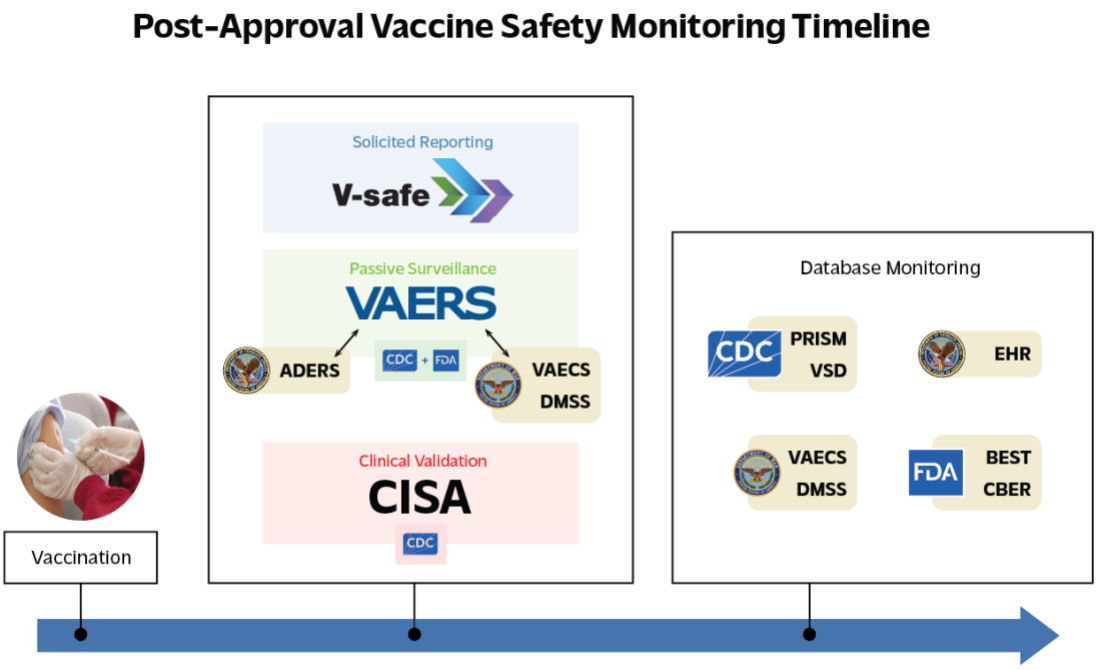

The governmental process involves both pre- and post-licensure and authorization scrutiny by both the CDC and FDA, which the CDC characterizes as "interagency collaboration to monitor post-licensure safety:"

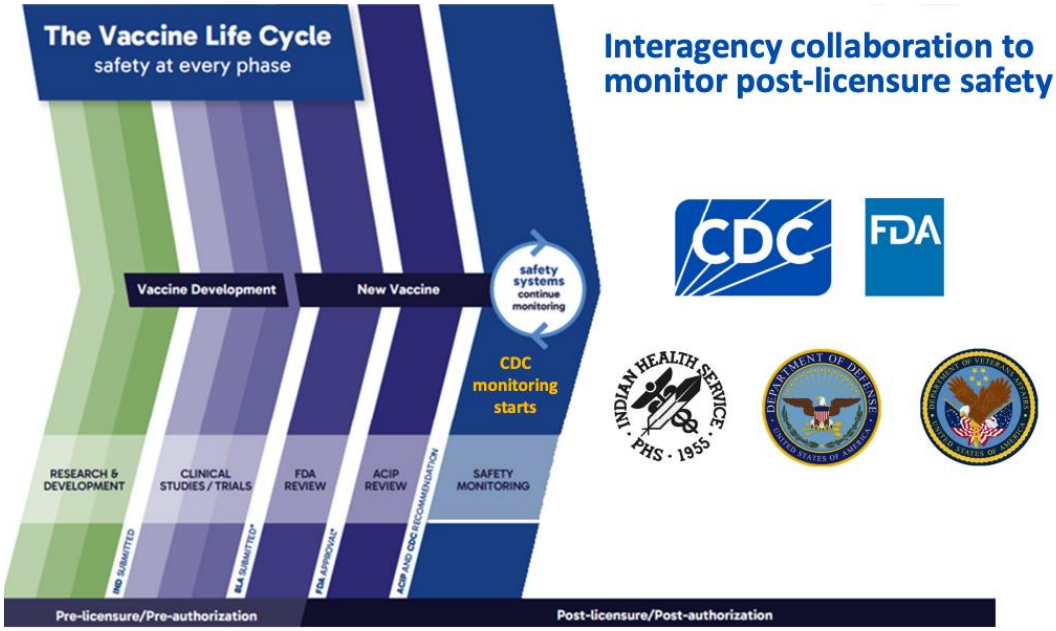

Sarah Meyer, MD MPH Director, CDC's Immunization Safety Office, "Update on CDC's COVID-19 Vaccine Safety Monitoring" (Meyer, 2025)

The process that led to CDC identification and validation of a causal link between an early COVID-19 vaccination and myocarditis in a specific subpopulation illustrates the range of resources in the federal vaccine safety monitoring system. CDC researchers report that data from VAERS, Vaccine Safety Datalink (VSD), the Biologics Effectiveness and Safety (BEST) System and the Clinical Immunization Safety Assessment Project (CISA), among other sources, led that agency to conclude that "the body of evidence from multiple U.S. and global monitoring systems was sufficient to establish a causal relationship between mRNA COVID-19 vaccination and myocarditis among adolescents and young adults, particularly among males and after dose 2"(Gee et al., 2024).

In June 2021, the FDA added a warning about myocarditis to mRNA vaccine labels.

U.S. Vaccine Safety Monitoring System: A Comprehensive, Coordinated, Complementary Structure

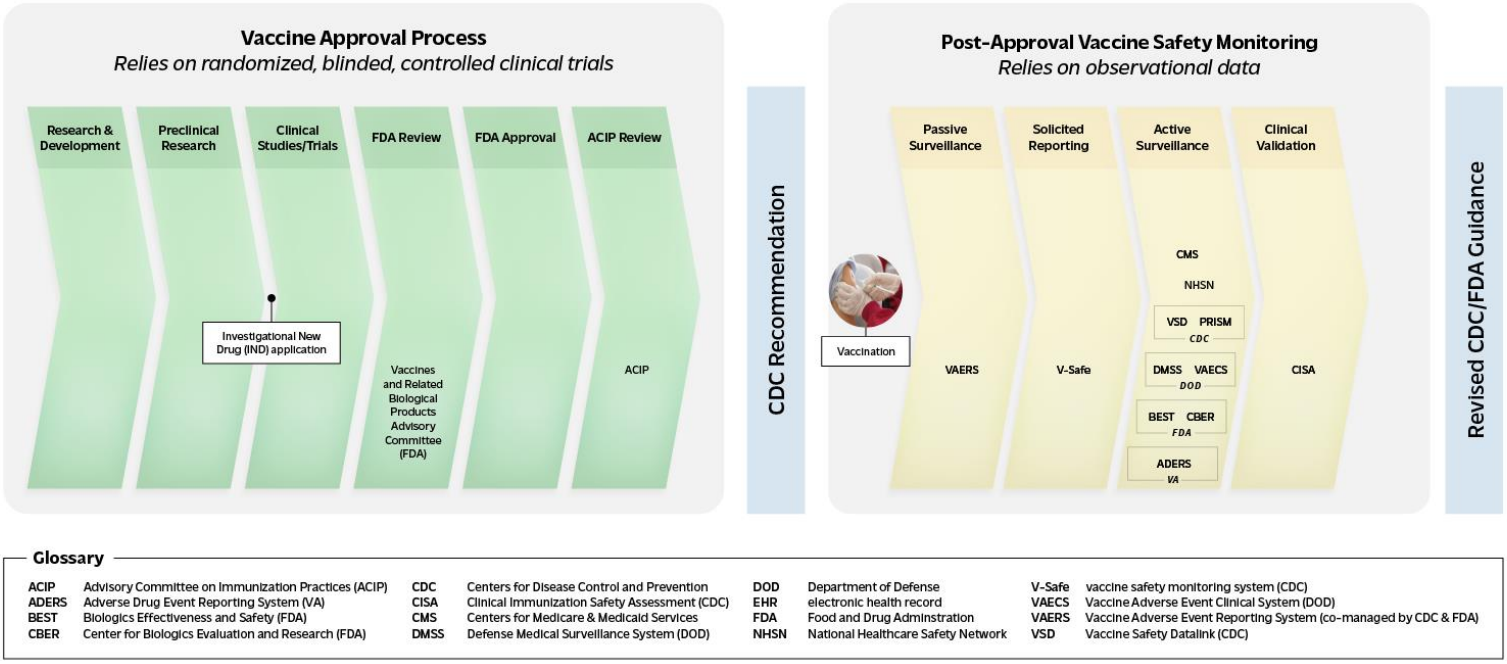

To increase understanding of the capacities of a number of the key entities in the federal safety monitoring system (indicated in the figure) and to make the case that they should be understood as parts of a coordinated, comprehensive vaccine safety monitoring system, Professor Jeffrey Morris, Ph.D., the George S. Pepper Professor of Public Health and Preventative Medicine at the Perelman School of Medicine of the University of Pennsylvania and a distinguished research fellow of the Annenberg Public Policy Center (APPC), authored the APPC white paper *The Complementary Components*

of the U.S. Vaccine Safety Monitoring System. (Professor Morris's paper is the third in the APPC vaccination communication toolkit. The first, issued in May 2023, was titled *Minimizing Public Susceptibility to Misconceptions About the Effects of Vaccination: Vaccine Adverse Event Reporting System*. The second, published in August 2023, focused on *Reducing Susceptibilities to Misconceptions About Vaccination During Pregnancy: RSV*.)

Professor Morris, who directs the Division of Biostatistics at the Perelman Medical School, is a fellow of the Institute for Mathematical Statistics and of the American Statistical Association (ASA) and a recipient of the ASA's Gottfried E. Noether Early Career Scholar Award. He also has held the Myrto Lefkopoulou Distinguished Lectureship at Harvard University and has won the M.D. Anderson Faculty Mentoring Award.

Kathleen Hall Jamieson, Ph.D.

Director, Annenberg Public Policy Center

University of Pennsylvania

December 2025

Abstract

Coming out of the COVID-19 pandemic, attention to vaccine risks has intensified. Detecting post-vaccination risks is inherently difficult, especially for events that are rare or occur long after vaccination. Randomized clinical trials can rigorously estimate common side effects and overall serious adverse events, but their limited size and duration constrain detection of rare or long-latency outcomes. Hence the FDA mandates post-approval safety monitoring.

The U.S. vaccine safety monitoring system is multicomponent and coordinated. Clinical case reports and passive surveillance (e.g., VAERS) generate hypotheses by flagging potential safety signals. Active electronic health record (EHR)-based systems, the CDC's VSD and the FDA's PRISM, are used to test and validate these signals to confirm links with vaccination, and then ideally are replicated in other large U.S. datasets (VA, DOD, CMS), in integrated health systems abroad, and via international networks (e.g., GVDN). The CISA clinical network complements this work with case-level review, mechanistic insight, and clinician guidance.

Despite heightened interest, public understanding of the system's architecture, and how its parts interact to form an integrated safety net, remains limited. This article reviews each component's role, strengths, limitations, and analytic considerations; explains how they are intended to function together; and proposes component-level and system-wide improvements to strengthen detection and evaluation of vaccine-related risks. The goal is better public explanation, constructive debate on enhancements, and a robust safety infrastructure that earns confidence.

Introduction

Vaccine safety has become an issue of renewed concern in the U.S. in recent decades. In the early 20th century, there was a fair amount of public skepticism and controversy regarding vaccines. However, the successful polio vaccine rollout in the mid-20th century marked a turning point, paving the way for routine immunization against other infectious diseases such as diphtheria, pertussis, measles, and rubella, among others. The rapid major decline that followed in illnesses, complications, and deaths from these diseases helped to convince much of the public that the health benefits far outweighed the potential risks.

Ironically, the success of these efforts also contributed to increasing societal concerns about vaccine safety. The resulting near-eradication of many vaccine-preventable diseases reduced public perception of their threat, decreasing the threshold of acceptable risk, and making even very small risks of a serious adverse event seem more significant by comparison. These concerns intensified during the worldwide COVID-19 pandemic, fueled by the rapid development and deployment of a new vaccine technology and widespread implementation of controversial vaccine mandates in many places.

These concerns brought increased attention to some of the vaccine safety monitoring systems established by the FDA and CDC to track and evaluate potential vaccine risks. However, this attention was only partial, often focusing on isolated aspects of the system, with limited awareness of its various components, the specific roles they serve, their individual limitations, and how they are designed to work together as an integrated framework.

For example, consider Google search trends for "VAERS" and "Vaccine Safety Datalink," two key components of the U.S. vaccine safety monitoring system, between 2019 and 2022 (Google Trends). Following the rollout of COVID-19 vaccines in January 2021, searches for "VAERS" surged nearly 100-fold, >20x higher than any point in Google search history going back to 2004. By contrast, searches for "Vaccine Safety Datalink" remained relatively low. As discussed later, **VAERS** (Vaccine Adverse Event Reporting System) is a passive surveillance system designed to detect potential safety signals, or "hypotheses of harm," based on reports from individuals and healthcare providers. Vaccine Safety Datalink (**VSD**), by contrast, is a much more rigorous active surveillance system that uses electronic health record (EHR) data to evaluate whether these safety signals correspond to actual increased risks after vaccination. Data from VAERS are highly susceptible to misinterpretation or deliberate misuse, which can lead to massively exaggerated perceptions of risk, especially when considered in isolation from data from other safety monitoring system components such as VSD. The risk of

distortion is amplified in periods of heightened public awareness and safety concern, such as during the rollout of COVID-19 vaccines during the pandemic. A comprehensive assessment from all components of the vaccine safety monitoring system reveals some risks from COVID-19 vaccines, but none approaching the magnitude commonly portrayed on social media through claims based on flawed or misleading use of VAERS data.

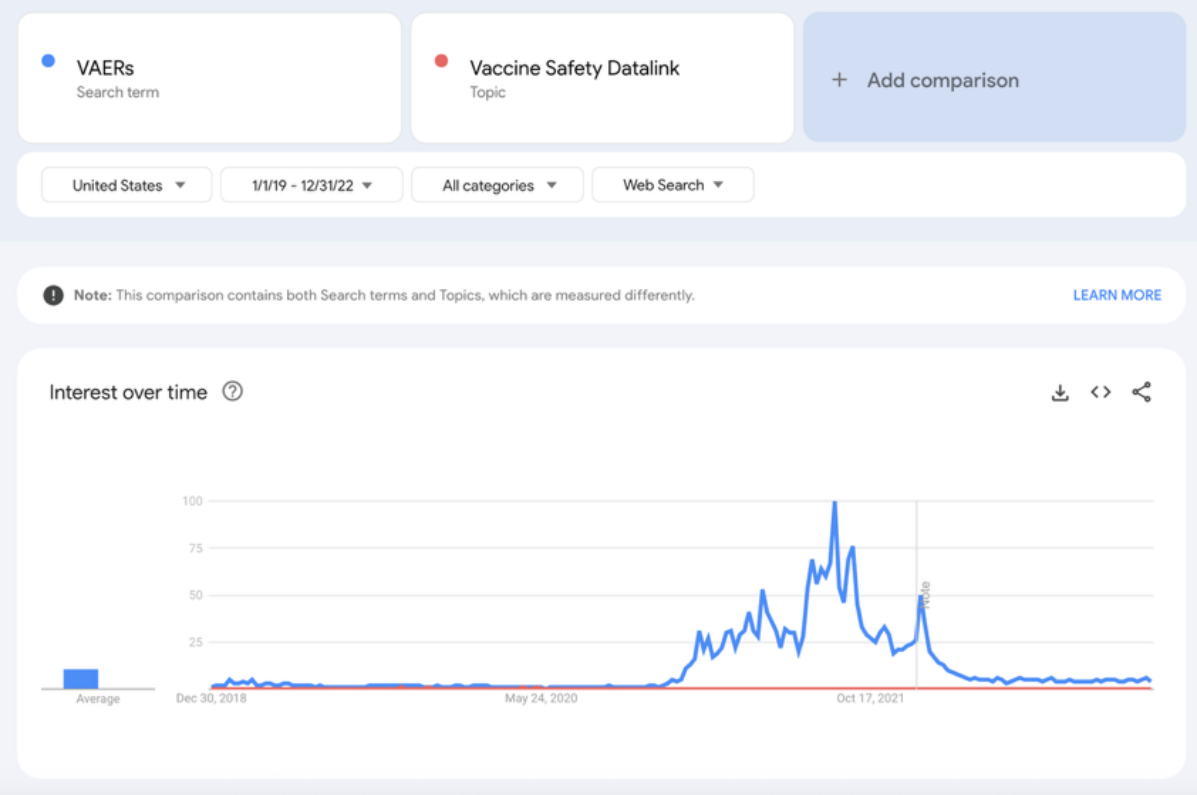

Since the early 1990s, the four major components of the vaccine safety assessment system include (1) **clinical trials** to demonstrate levels of efficacy and safety required for approval; (2) **VAERS**, a passive surveillance system to detect potential safety signals; (3) the CDC's **VSD,** FDA's **PRISM** monitoring program within the **Sentinel system,** and other government medical records systems used to further validate safety signals detected by VAERS, including other government data sources from the Veteran's Administration (VA), the Department of Defense (DOD), and the Center for Medicare & Medicaid Services (CMS); and (4) the Clinical Immunization Safety Assessment **(CISA**), a national network of vaccine safety experts able to perform clinical case consultations for individual patients experiencing post-vaccine adverse events. During the COVID-19 pandemic, the new system **V-safe** was developed to gather real-time safety monitoring information for COVID-19 vaccines from short check-ins on mobile devices and was

expanded for use with RSV (respiratory syncytial virus) vaccines, with future expansion possible.

In this article, I will provide an overview of the key components of the U.S. vaccine safety monitoring system, outlining the type of data each component generates, their respective strengths and limitations, and important considerations for their analysis and interpretation. I will clarify the specific role of each component and explain how they are designed to work together to provide a more comprehensive view of vaccine safety. A full understanding of potential vaccine safety concerns requires integrating information from across the entire system. While acknowledging the system's limitations, I will also suggest various ways it could be strengthened and improved. Although the discussion is broadly applicable to vaccine safety monitoring as a whole, I will place particular emphasis on the challenges and lessons that emerged during the COVID-19 pandemic.

Clinical Trials

The first source of safety information for a new vaccine comes from the pre-licensure clinical trials required by the FDA for approval. These are typically randomized controlled trials (RCT) that compare individuals receiving the vaccine to those in a control group, who are given either a placebo or another vaccine. RCTs evaluate both efficacy, such as reduction of infection risk as well as immunogenicity, and safety by comparing adverse event rates between the vaccinated and control groups. Trial sizes typically range from several hundred to tens of thousands of participants per group.

Unlike observational studies, where the decision to vaccinate may be influenced by factors that induce bias, RCTs use random assignment to eliminate such biases in treatment allocation. This leads to unbiased estimation of vaccine effects, one reason why RCTs are considered the "gold standard" in scientific research, especially double-blind trials in which neither participants nor clinicians know the treatment assignment during the blinding period.

Safety monitoring in vaccine trials is done by recording all reported adverse events (AEs) and summarizing their incidence in vaccinated and control groups through a series of tables and figures in the clinical trial report. The AEs are tracked over various time frames: short (0-7 days), intermediate (<30 days), and longer term (through 6 months or more). The short-term AEs include "solicited" AEs that are predefined and expected reactions to immune stimulation effects that could be local (e.g., redness or pain at the injection site) or systemic (e.g., fever or fatigue) and are actively monitored.

The intermediate AEs include "unsolicited" AEs, which are not specifically solicited but reported by participants or observed by investigators.

Longer-term AEs include **serious adverse events** (SAE) and **adverse events of special interest** (AESI). SAEs are events that result in hospitalization, require medical intervention, are life-threatening, cause significant disability, or lead to death. They are distinct from **severe adverse events**, which refer to the intensity of an event rather than its medical seriousness. For example, an intense headache that prevents someone from going to work the next day may be classified as a *severe* event, but it would not be considered *serious* unless it meets criteria such as requiring medical attention or hospitalization. AESIs are predefined, theoretically plausible risks based on prior knowledge of the vaccine or vaccine platform, e.g., myocarditis, anaphylaxis, or Guillain-Barré Syndrome (GBS), that are actively monitored, even if rare, due to their potential seriousness. Follow-up for SAEs and AESIs typically continues for at least six months after vaccination, and in some trials may extend longer, depending on the study design and regulatory requirements.

All AEs are required to be reported in clinical trials regardless of whether they are believed to be caused by the vaccines. Therefore, it is essential to compare AE rates between the vaccinated and control groups to determine whether any reported events occur more frequently following the vaccine under investigation, which would suggest a potential causal relationship. The randomized trial design provides a rigorous framework for statistically evaluating adverse event incidence between groups.

While many clinical trials are adequately powered to detect differences in the most common adverse events or aggregate safety measures, for example, the overall rate of SAEs, they are not large enough to reliably detect differences in the rates of **rare individual adverse events**. For example, a trial with 20,000 participants per group would have sufficient power, 80% power with 0.05 type 1 error level, to detect a modest increase (e.g., from 1 in 60 to 1 in 50) in overall SAE rates (or 1.2-fold increase). However, detecting differences in rare individual events requires much larger sample sizes:

- For events with population incidence of 1 in 1,000, a trial with 20,000 per arm would have sufficient power to detect a 2.2-fold increase (to about 1 in 500) in the vaccination group.
- For events with population incidence of 1 in 10,000, the same trial would only be powered to detect more than a 7.2-fold increase (to about 1 in 1,500).
- For events with a population incidence of 1 in 100,000, it would only be powered to detect increases greater than 50-fold (to about 1 in 2,000).

To detect smaller relative increases in such rare events, a trial would have to be extremely large. For example, detecting a 2-fold increase in an event with control

incidence of 1 in 10,000 would require >250,000 participants per group, and detecting a 2-fold increase in an event with control incidence of 1 in 100,000 would require more than 2.5 million participants per group. Sample sizes such as these would be too large to be considered feasible.

This represents a key limitation of the RCTs when it comes to safety monitoring: although their rigorous design enables unbiased estimation of vaccine effects, **their sample sizes are too small to detect moderate to rare adverse events**, e.g., those with incidences on the order of 1 in 10,000 to 1 in 100,000. Such event rates, while statistically rare, may still be of significant concern for a product like a vaccine that is administered to large, generally healthy populations.

RCTs also have additional limitations when it comes to detecting real-world post-vaccination risks, particularly regarding their **representativeness across population, place, and time**. First, clinical trial participants are often not representative of the general population. They tend to be younger and healthier and may exclude or underrepresent vulnerable groups such as the elderly, immunocompromised, or those with serious comorbidities. Second, trials are frequently conducted in wealthier regions and controlled environments such as academic medical centers, raising questions about how well the findings generalize to other geographic or healthcare settings. Third, RCTs are conducted over a fixed time period and thus do not capture how vaccine performance or risks may evolve under changing real-world conditions. During the COVID-19 pandemic, for example, these conditions included major shifts in societal mitigation efforts, the emergence of new viral variants, changes in past infection history, and rising levels of population immunity, all of which could influence vaccine safety and effectiveness measures and their relative significance.

A further limitation is the challenge of assessing potential **long-term sequelae** years after vaccination. Sustaining blinding and randomization over extended periods is rarely feasible and raises ethical, logistical, and accrual issues; even when attempted, trials face crossover, off-protocol vaccination, and unblinding. Consequently, RCTs are well suited to detecting near-term safety signals but have inherent limitations for rare or long-latency outcomes.

These limitations highlight the importance of robust **post-licensure studies** and **post-approval safety monitoring** with the potential to detect rare or longer-latency serious adverse events after vaccination in real-world settings to serve as an essential complement to the pre-licensure randomized clinical trials. The current vaccine safety monitoring framework in the United States includes a set of complementary systems, each designed to address different aspects of post-market surveillance, which we will now examine in detail.

VAERS, a passive surveillance system

One outcome of the National Childhood Vaccine Injury Act (NCVIA) in 1986 was the establishment of a unified national system to detect and assess rare adverse events after vaccination. This included the creation in 1990 of the Vaccine Adverse Event Reporting System (VAERS), co-administered by the Centers for Disease Control and Prevention (CDC) and the Food and Drug Administration (FDA). VAERS is a passive surveillance system (Shimabukuro et al., 2015) designed to collect reports of adverse events from vaccine manufacturers, healthcare providers, and members of the public. Its purpose is to help identify rare or unexpected events that may not have been detectable in pre-licensure clinical trials.

Because it is passive, VAERS relies primarily on voluntary submissions, with no active follow-up or systematic data collection. Reports can be submitted by individuals who experienced the event, their family members, healthcare professionals, or manufacturers. Reporting is encouraged even if the reporter is not sure whether the vaccine caused the event; healthcare workers are required by law to report certain events (VAERS - Report an Adverse Event), and manufacturers are required by law to submit adverse events of which they become aware. Reports submitted to VAERS can be reviewed for completeness, and a basic coding of the events is done using medical coding systems such as **MedDRA** (Medical Dictionary for Regulatory Activities) or **OAE** (Ontology of Adverse Events). VAERS staff will follow up on some reports deemed serious by contacting healthcare providers for medical records, death certificates, or autopsy reports. However, not all serious events are independently verified, and nonserious ones are rarely followed up on, particularly during periods with high report volume, such as during the COVID-19 pandemic. Importantly, VAERS is not designed to assess causality, but instead to serve as an early warning system for detecting potential **safety signals** that warrant further investigation.

The strengths of VAERS include its national scope and its ability to potentially pick up on very rare events, which makes it well suited for its primary purpose of hypothesis generation. Another important strength is its transparency; the data are publicly accessible, allowing independent researchers and the public to examine reported adverse events directly.

However, passive surveillance systems like VAERS also have several important limitations that must be carefully considered to ensure accurate interpretation of their data, including:

1. Variability in quality and completeness of reports
2. Reporting bias of variable and unknown magnitude

3. Lack of an unvaccinated control group

I will now discuss each of these limitations.

One key limitation is the **variability in the quality and completeness of reports**, which often lack sufficient clinical detail to confirm the medical diagnosis or to enable meaningful follow-up by VAERS staff. Reported events are often not medically verified, and in some cases are completely uncorroborated.

Another key limitation is **reporting bias of unknown magnitude,** with many events occurring after vaccination going underreported, and the degree of underreporting varying widely across types of event and the context in which they occur. Reporting is more frequent for serious events compared to mild ones (Miller et al., 2020; Rosenthal & Chen, 1995; McNeil et al., 2013), and for events that occur closer in time to vaccination than those occurring later (Rosenthal & Chen, 1995; Braun & Ellenberg, 1997). Additionally, **stimulated reporting,** increased reporting due to media coverage or social media attention, can lead to overreporting of certain events relative to others (Miller et al., 2020; Eberth et al., 2014; Verstraeten et al., 2001). This phenomenon is particularly relevant for specific widely publicized adverse events, newly introduced vaccines (Verstraeten et al., 2001), or exceptional circumstances such as the COVID-19 pandemic. One study found that nearly half of all VAERS reports filed since its inception in 1990 were filed in 2021 (Almadani & Alshammari, 2022) and, as noted earlier, Google trends data show a dramatic spike in public awareness of VAERS during that same year (Google Trends). Because of this reporting bias and the passive nature of the system, it is not possible to accurately calculate incidence rates for adverse events using VAERS data, as both the numerator (number of true cases) and denominator (number of individuals at risk) are not reliably known.

A third key limitation of VAERS is the absence of an **unvaccinated comparison group**. Since all adverse events have a background incidence in the population, some events will occur by chance following vaccination and may be reported to VAERS regardless of any causal relationship. As a result, one cannot assume that an event was caused by vaccination simply because it was reported. Without a suitable control group, it is difficult to determine whether a particular event occurs more frequently after vaccination or whether reported cases simply reflect coincidental background occurrences unrelated to the vaccine. This limitation also makes it impossible to accurately estimate absolute or relative risks using VAERS data alone.

These limitations render VAERS unsuitable for determining causal relationships between vaccines and adverse events. Instead, it is designed to serve as an early warning system for detecting **potential safety signals** or **hypotheses of harm** that warrant further investigation through more rigorous studies that are less susceptible to these limitations. VAERS analyses are primarily designed to detect population-level

safety signals, which is the focus of this section. However, as I will discuss later in the section on the CISA system, detailed review of individual VAERS reports can also produce case studies that uncover important insights and help generate new hypotheses. Below, I will outline several statistical methods commonly employed by researchers and public health officials to identify potential safety signals in VAERS data while accounting for these constraints.

However, because VAERS data are publicly accessible, some individuals conduct their own analysis using flawed or inappropriate methods that fail to account for these inherent limitations, which can result in misinterpretation or misrepresentations of the data, often leading to highly exaggerated claims of vaccine risk. Before describing legitimate methods to analyze VAERS data, I will first describe four common invalid analytical approaches I have frequently encountered in claims made by individuals on social media and in mainstream media, including public officials and other influential figures, and explain why they are flawed. They are:

1. Presenting raw VAERs report counts as confirmed vaccine-caused events.
2. **Multiplying report counts by an assumed underreporting factor** and treating the result as an estimate of vaccine-caused events.
3. Inferring vaccine causality from increases in raw VAERs report numbers across different vaccines or over time.
4. Claiming causality based on temporality, simply because **adverse event reports decline with days since vaccination**.

First, some individuals **present raw numbers of VAERS reports as though they represent confirmed vaccine-caused events**, disregarding the fact that VAERS does not establish causality. This misinterpretation also overlooks the reality that a certain number of events are expected to occur by chance alone, based on their background incidence in the general population, even in the absence of vaccination. This issue is particularly concerning for the most serious events, which healthcare providers are legally required to report to VAERS, regardless of whether they are believed to be related to vaccination.

Second, this misinterpretation is sometimes further compounded by multiplying the raw number of VAERS reports by an assumed **underreporting factor (*URF*)** and presenting the resulting figure as an estimate of vaccine-caused events, with the *URF* defined as:

$$\boldsymbol{URF} = \text{true number of events occurring in the population/events reported to VAERs}$$

This approach is fundamentally flawed, not only because it misrepresents causality and ignores background event rates, but also because it assumes a uniform underreporting factor across events and conditions, which is inaccurate. Many who use this method assume *URF* = 100, citing a statement that "fewer than 1% of vaccine adverse events are reported" that appeared in a 010 report prepared for the Agency for Healthcare Research and Quality (AHRQ) (Lazarus, 2010). However, this figure did not come from original data collected in that report. It was a general comment made by the authors to underscore the value of automated reporting systems over passive surveillance, and it included mild events such as injection site redness and soreness, which are known to be severely underreported relative to more serious events.

As discussed above, reporting sensitivity varies considerably across event types and contexts. Serious events are more likely to be reported, particularly when they occur soon after vaccination or during periods of heightened public attention such as the COVID-19 pandemic. Empirical studies have found much higher reporting rates for specific serious events. For example, one study (Miller et al., 2020) comparing VAERS with data from the VSD estimated that VAERS captured 13-76% of Guillain-Barré Syndrome (GBS) cases and 13-27% of anaphylaxis cases, depending on the vaccine. Another study using a capture-recapture method that linked VAERS data to case-control and retrospective cohort studies found a 47% reporting sensitivity for intussusception after rotavirus vaccination (Verstraeten et al., 2001).

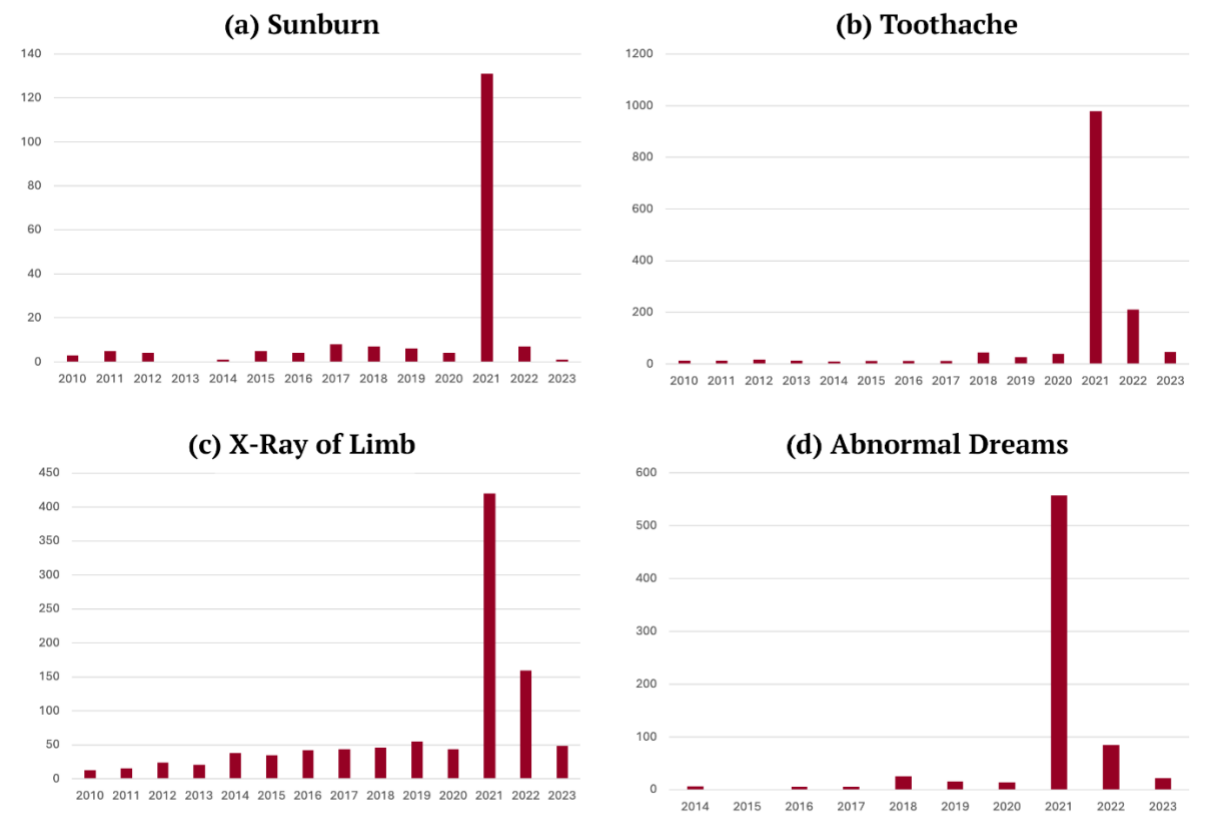

Third, another flawed approach is to **infer vaccine causality from increases in reported adverse events across vaccines or over time**. This method fails to account for various key factors: some vaccines were administered to far more people than others or to different subgroups with varying baseline risks, reporting rates may vary across vaccines, and reporting patterns may fluctuate considerably depending on the timing and context. This issue became particularly pronounced during the COVID-19 pandemic, when individuals compared VAERS report counts for COVID-19 vaccines in 2021 to reports for other vaccines from prior years, and took any sharp increase in 2021 to imply a causal vaccine effect. This approach is still regularly seen in claims circulating on social media. As previously noted, there was a major spike in VAERS awareness and reporting in 2021 (Google - Google Trends; Almadani & Alshammari, 2022). In fact, nearly every type of event reported to VAERS, including many with no obvious link to vaccines such as sunburn, toothache, X-ray of limbs, and abnormal dreams, showed dramatic spikes in 2021, often increasing by orders of magnitude (The Vaccine Adverse Event Reporting System (VAERS) Request) (see figure). This widespread surge makes clear that a sharp rise in VAERS reports for COVID-19 vaccines in 2021 cannot, on its own, be taken as evidence of causality. Instead, it reflects a general elevation in reporting rates across the board, driven by heightened awareness, intense public scrutiny, and unprecedented vaccine rollout.

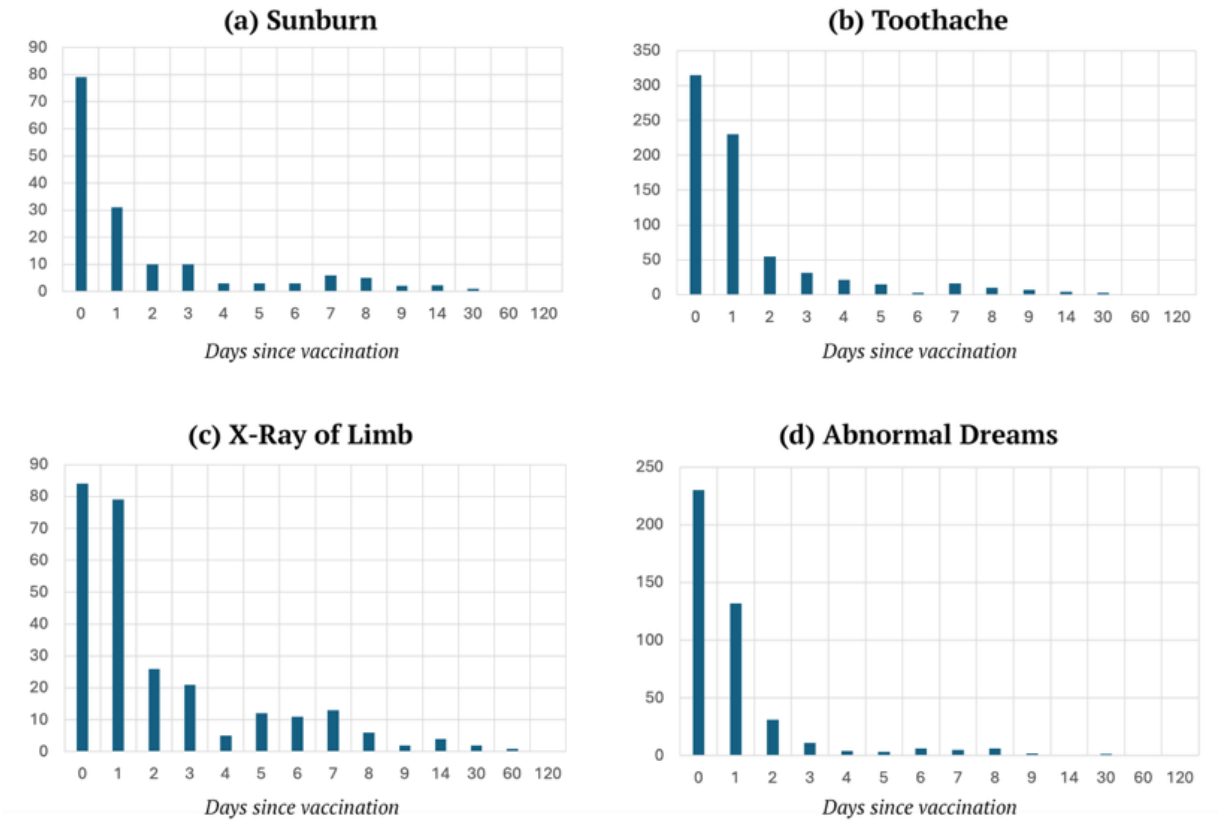

Fourth, citing the fourth of Bradford Hill's conditions (Hill, 1965) for causality, *temporality*, some have argued that a decline in adverse event reports over time since vaccination indicates a causal link. However, this reflects a misapplication of the temporality criterion, as it ignores the well-documented tendency for all adverse events, regardless of link to vaccination, to be reported more frequently shortly after vaccination (Scholl et al., 2025). Braun and Ellenberg (1994) showed that 45% of events reported to VAERS occurred on the day of vaccination, an additional 20% on the following day, and a total of 92.5% within two weeks. VAERS data consistently show that reports for a wide variety of events, including those with no obvious link to vaccination such as sunburn, toothache, X-ray of limb, and abnormal dreams, are highly concentrated in the days immediately following vaccination and decline rapidly thereafter (see figure). This pattern is an inherent feature of passive reporting systems and does not in itself indicate causality.

Not only are the causal claims that often accompany these four simplistic methods unfounded, but the methods themselves are fundamentally flawed and poorly suited even for hypothesis generation of vaccine-specific safety signals. In contrast, researchers and public health officials rely on a range of more robust statistical approaches designed to identify potential safety signals in VAERS data that better account for their inherent limitations. These include:

1. Comparing reporting rate with historical background rate
2. Proportional Reporting Ratios (PRR)
3. Empirical Bayes approaches (EB)

I will briefly describe each of these approaches, highlighting their underlying assumptions and key limitations, and suggest some potential improvements.

One commonly used approach involves calculating the **VAERS reporting rate** for a specific adverse event and then comparing it to the **historical background rate** of that event in the population. This comparison can help assess whether the event is reported more frequently following vaccination than would be expected by chance. As previously noted, VAERS data cannot be used to estimate true incidence rates due to uncertainty in both the numerator (number of true cases) and denominator (number of individuals at risk). However, one can estimate a **VAERS reporting rate** by dividing the number of VAERs reports for an event by the estimated number of vaccine doses regionally administered during the corresponding time period. This rate can then be compared to the background incidence rate, derived from historical data, to determine whether the frequency of reports exceeds what would typically be expected, suggesting a potential signal.

However, this approach as described has important flaws and limitations. Chief among them is that it does not account for **underreporting**, a well-known feature of passive surveillance systems like VAERS. While the true underreporting factor (*URF)*, defined as the ratio of actual to reported event counts, is unknown and can vary substantially across different types of events, underreporting is clearly a real and influential factor and should not be assumed to be negligible. As discussed earlier, relying on a fixed *URF* without strong empirical justification is problematic, reinforcing the need for a more nuanced and flexible approach to incorporate the variability and uncertainty in the degree of underreporting.

One potential improvement to this approach is to explore the implications of underreporting by examining a range of plausible *URF*s. To formalize this, let the ***VB* ratio** represent the ratio of VAERS reporting rate to the historical background rate. If the true *URF* were known, the **relative risk (*RR*)** of the adverse event following vaccination, relative to historical background, could be approximated by:

$$RR = VB \times URF$$

A value of $RR >> 1$ would indicate a potential safety signal. Although the *URF* is unknown, one can calculate the **minimum URF** required for $RR > 1$, which is simply

1/VB. This provides a form of sensitivity analysis, assessing how large the underreporting factor would need to be for the observed data to suggest elevated risk.

When using this approach, it is essential to obtain a reliable estimate of the historical background incidence rate, ideally from a population-level model that can be calibrated to reflect the demographic and clinical characteristics of the vaccinated population. Additionally, one must consider that the historical background rates may not be applicable during periods of rapid change. For example, during the COVID-19 pandemic, factors such as SARS-CoV-2 infections, widespread mitigation efforts, altered diet and activity patterns, and reduced healthcare access may have independently altered the risk of certain adverse events, potentially rendering historical baselines inaccurate for valid comparison.

Given a formal model for the historical background incidence rate for the vaccinated subpopulation and specifying a prior distribution reflecting plausible values of the underreporting factor (*URF*), one could perform **fully Bayesian analysis** to estimate the relative risk (RR) and calculate the posterior probability that $RR > 1$, representing the probability that the modeled event constitutes a safety signal. Given the inherent uncertainty and variability of the *URF*, such an approach would require a **careful sensitivity analysis** on the choice of prior distribution for the *URF* to ensure robustness of the conclusions. To date, I am not aware of any published efforts that have implemented such a fully Bayesian modeling approach in the vaccine safety surveillance literature, but to me it is a compelling idea.

A second and the most common analytical approach used to identify potential safety signals is the **proportional reporting ratio (*PRR*),** which assesses whether the ***relative*** reporting rate of a given adverse event (AE) of interest is greater for the vaccine of interest than other vaccines. This approach was first introduced by Finney (1974) and popularized by a paper by Evans et al. (2001). The *PRR* is computed from the following 2×2 table:

	Reports for AE of interest	Reports for other AEs
Vaccine of interest	A	B
Other vaccine(s)	C	D

(1)

From the data contained in Table (1), the ***PRR*** is computed as:

$$PRR = \frac{A/(A+B)}{C/(C+D)}$$

Large *PRR* indicate that the relative reporting rate for the AE of interest is greater for the specified vaccine than other vaccines, demonstrating an association between the specified AE and vaccine of interest, and suggesting a potential safety signal. Statistical significance can be assessed using confidence intervals or χ^2 tests. The ***screened proportional reporting ratio*** (***SPRR***) (Evans et al., 2001) is one commonly used criterion that flags an AE of interest as a potential safety signal if *PRR*>2, χ^2 >4, and A>3.

Unlike the simplistic comparisons of raw report numbers described earlier, the *PRR* adjusts for differences in vaccination and overall VAERS reporting rates across vaccines and contexts, enabling more meaningful comparisons of reporting patterns between vaccines or across time periods. However, the method also has important limitations that must be considered.

First, it may miss safety signals if a vaccine increases the risk of a broad array of adverse events, thereby diluting the signal for any single event. Second, it cannot account for event-specific increases in reporting driven by factors unrelated to the vaccine itself, such as stimulated reporting that disproportionately affects one vaccine over others or time-related shifts in the underlying risk of the event during the vaccination period relative to the time period used to estimate the background. Third, if the vaccinated populations differ substantially in baseline event risk, this could lead to false positives or false negatives in safety signal detection. For example, one would naturally expect more reports of myocardial infarction after shingles vaccines, given primarily to older adults, than MMR vaccine, given primarily to children. All of these issues are particularly relevant in the context of the COVID-19 pandemic. As a result, the *PRR* approach is vulnerable to both missed signals and spurious associations.

Sarntivijai et al. (2012) discussed an alternative application of the Proportional Reporting Ratio (*PRR*) that uses the total number of regional vaccine doses administered rather than the total number of reports for all events as the normalization factor. This approach adjusts for differences in vaccination rates across regions but does not account for broader variations in VAERS reporting propensity across vaccines or time periods, so is less robust, especially in contexts with known stimulated reporting, such as during the COVID-19 pandemic.

A third commonly used approach for detecting VAERS safety signals is to use an **Empirical Bayes (EB)** model (DuMouchel 1999; Banks et al., 2005). This approach builds on the same foundational concepts as the *PRR* but incorporates a Bayesian framework to better account for sampling variability and to compute probabilities that help determine whether the observed signal is strong enough to warrant concern.

The core idea is to model the number of reports for the AE of interest for the vaccine of interest (denoted as *A*) using a Poisson distribution with mean parameter μ. The goal is

then to assess the evidence that μ exceeds a baseline expectation E, which represents the expected number of reports under the assumption that reporting of the AE is independent of vaccine type. The expectation is calculated as:

$$E = \frac{(A + B)(A + C)}{A + B + C + D}$$

where A, B, C, and D are defined as in Table (1) above. The ratio $\lambda = \mu/E > 1$ quantifies the deviation from expectation, and the EB method evaluates the posterior probability that $\lambda > 1$, that is, the probability that the observed reporting rate exceeds what would be expected under the null hypothesis of no association.

Technical details include specification of a prior distribution for λ that is a mixture of two Gamma distributions. One component is tightly centered around 1, representing the null hypothesis of no association, while the other has greater variability and allows for deviations suggestive of a potential safety signal. The hyperparameters of these distributions are estimated directly from the data, which is what makes the method **empirical Bayes**. This prior structure introduces **shrinkage** for small values of A, helping to stabilize estimates and more accurately reflect uncertainty in low-count scenarios.

This method produces a statistic called the **Empirical Bayes Geometric Mean (*EBGM*)**

$$EBGM = E(\log_2 \lambda | A),$$

where $E(X|A)$ denotes the expected value of a random variable X while conditioning on observed data A. Exponentiating yields an estimate of the signal ratio $2^{EBGM} \approx \lambda$. In addition, the method computes a **posterior probability**:

$$P = \Pr(\lambda > 1 | A)$$

which quantifies the probability that the adverse event (AE) is reported more frequently than expected under the null and can be interpreted as a probability of being a safety signal. A commonly used threshold for flagging potential safety signals is $P > 0.95$, a criterion sometimes referred to as ***EB05.*** (Banks et al., 2005; DuMouchel & Pregibon, 2001).

This method shares many of the same strengths and limitations as the *PRR*. It accounts for differences in vaccination rates and overall VAERS reporting rates across vaccines and contexts but may still miss safety signals if a vaccine increases the risk of a broad spectrum of adverse events, thereby diluting the signal for any one event. Similarly, it may flag false safety signals in situations in which event-specific increases in reporting are driven by external factors unrelated to the vaccine itself or if the baseline risk differs across the vaccinated populations compared. As with the *PRR*, it is possible to calculate the baseline expectation *E* using the total number of regional vaccine doses administered instead of the total number of reports for all events. However, as above, this adjustment would only account for differences in vaccination rates but not address potential biases in overall VAERS reporting rates across vaccines or time periods, and so would not be as robust.

Key advantages of the Empirical Bayes modeling approach include its use of shrinkage for low-count data and its ability to compute posterior probabilities of association, which provide a more interpretable and nuanced assessment of whether an adverse event may represent a true safety signal. The *EB05* criterion, which flags signals when the posterior probability that $\lambda > 1$ exceeds 95%, helps reduce false positives by applying a stringent evidence threshold. However, this conservatism comes at the cost of increased false negatives, potentially missing some true safety signals that may be detected by other methods (Banks et al., 2005).

While these methods are valid tools for detecting potential safety signals, they are still inherently limited by the constraints of a passive surveillance system. As such, signals identified through these analyses should not be interpreted as definitive evidence of vaccine-related risk, but rather as **hypotheses of harm** that warrant further investigation and validation, for example using other components of the vaccine safety monitoring system such as active reporting systems.

One potential improvement to the current system would be the integration of automated adverse event reporting within healthcare settings (Lazarus, 2010). This could enhance surveillance efforts by reducing underreporting and minimizing reporting bias. However, even with such advancements, the lack of an appropriate, contemporaneous unvaccinated control group would still hinder the system's ability to assess whether certain events are occurring more frequently following the vaccine in question. This highlights the continued need for follow-up analyses that compare reporting rates to well-chosen control groups, an approach routinely used in active surveillance systems such as the Vaccine Safety Datalink (VSD) to evaluate potential safety signals.

Vaccine Safety Datalink (*VSD*), an active surveillance system

Another key outcome of the National Childhood Vaccine Injury Act (NCVIA) of 1986 was the creation of the **Vaccine Safety Datalink (*VSD*)**, an **active surveillance system** designed to detect and validate rare vaccine-associated risks (Chen et al., 1997). Launched in 1990, the VSD originally established a network of four managed care organizations (MCOs) integrating their electronic health records (EHR) using a common data model and standardized coding system. Since its inception, the VSD has expanded to include EHRs from nine MCOs (CDC 2025; McNeil et al., 2014), covering >10 million members (Klein et al., 2021), or more than 3% of the U.S. population. The *VSD* population has been shown to be reasonably representative of demographics of the broader U.S. population (Sukumaran et al., 2015). This system is used to perform weekly **rapid cycle analysis (RCA)** for monitoring prespecified adverse events of special interest (AESI) and to facilitate ad hoc studies in response to potential safety signals identified through VAERs or other sources, whether initiated by researchers within or outside the VSD network.

As an active surveillance system, VSD does not rely on voluntary reporting of adverse events from healthcare providers or individuals. Instead, it utilizes electronic health records (EHRs) to systematically capture all adverse events medically verified and documented for both vaccinated and unvaccinated individuals using standardized ICD-9 and ICD-10 diagnostic codes. Data from the VSD are not subject to the same types of underreporting and reporting bias as passive surveillance systems such as VAERS, since the data are drawn directly from routine clinical care records rather than self-initiated reports.

However, as with any EHR-based system, the VSD only captures adverse events that lead to medical attention. Minor events that do not prompt a clinical visit are likely to go unrecorded. Additionally, adverse events in groups with lower healthcare utilization may be underrepresented, introducing **ascertainment bias**. This is especially true for outcomes prone to **under-ascertainment**, and for which the likelihood of diagnosis varies across subpopulations. The same factors that reduce the likelihood the condition is diagnosed and recorded in the medical records may also correlate with the propensity to be vaccinated. This concern is particularly salient when studying potential links

between early childhood vaccination and chronic conditions, many prone to under-ascertainment, that are often not diagnosed until years later.

Despite these limitations, most serious adverse events are documented in the electronic medical records, making the VSD a valuable and reliable resource for monitoring and evaluating potential vaccine-related safety concerns.

To assess whether the incidence rate of a given adverse event after vaccination suggests an increased risk, it is necessary to compare it with some suitable control group. Three key types of controls commonly used in VSD studies include:

1. **Unvaccinated controls**, comparing with unvaccinated individuals in the population.
2. **Historical controls**, comparing with historical background rate for the event.
3. **Risk-interval based controls**, comparing incidence in a pre-specified risk interval relative to a control interval, using either **self** or **concurrent comparator** controls.

Unvaccinated controls may seem like the most straightforward control group, allowing for direct comparison of adverse event rates between vaccinated and unvaccinated populations. However, this approach is often problematic and, in many cases, infeasible within the VSD. There is high vaccine coverage for many vaccines within the VSD population, and this often leaves too few unvaccinated individuals to serve as meaningful controls (Barosa & Prasad, 2025). Moreover, those who remain unvaccinated may be such a highly selective subgroup, differing substantially from vaccinated individuals in demographic, behavioral, health-related, and healthcare-seeking characteristics, that they make a poor comparison group.

In some cases, and for certain vaccines, there may be a sufficient number of unvaccinated individuals to enable meaningful comparative analyses. A notable example is the early phase of the COVID-19 vaccine rollout during the pandemic. However, because vaccination is not randomly assigned, comparisons between vaccinated and unvaccinated groups must account for systematic differences that could influence the risk of the outcome (known as **confounders**). These may include differences in baseline demographics, medical history, healthcare-seeking behavior, and relevant time-related factors, especially if the observation periods differ substantially between groups. For example, given that individuals who were elderly or at high risk were among the first to be given access to COVID-19 vaccines, reflecting **confounding by indication**, one would expect to see more coincidental adverse events among the early vaccinees than the younger cohorts awaiting their opportunity to get vaccinated.

One promising approach in these settings would be the use of a **trial emulation design** (Hernán et al., 2022), which aims to mimic the structure of a randomized trial within an

observational data context. This approach involves defining fixed enrollment and observation time windows, establishing clear inclusion criteria, and creating matched pairs of vaccinated and unvaccinated who are similar in baseline demographic and clinical characteristics. These pairs are then followed prospectively for the outcome of interest, and with follow-up ending for both individuals if the unvaccinated individual receives the vaccine, either individual experiences the outcome, or either is lost to follow-up or dies. Matching on calendar time, baseline factors, and enforcing identical follow-up windows reduces the biases of naïve vaccinated–unvaccinated comparisons and produces more credible effect estimates.

To enhance causal inference, the design can incorporate **propensity score reweighting** (Austin 2011) or **stratification** to adjust for additional potential confounders beyond those used in the initial matching. Furthermore, incorporating **negative control outcomes** (Lipsitch et al., 2010) by replicating the study design for outcomes that are expected to share similar confounding structures but not plausibly related to vaccination, can help assess the degree of residual confounding and strengthen the credibility of the findings. **E-values** (VanderWeele & Ding, 2017) can be used to estimate the minimum strength of association that an unmeasured confounder would need with both treatment and outcome to account for the observed effect, thereby assessing the robustness of results to residual confounding. Although not widely applied historically, trial emulation designs are an innovative cutting-edge approach that holds considerable promise in this context.

An alternative to a trial emulation design is a **cohort study treating vaccination status as a time-varying exposure**. This approach enables the evaluation of vaccine effects on the risk of a specified adverse event while accounting for changes in vaccination status over time. It is particularly useful when overall vaccine coverage is high but there is enough variability in vaccination timing to allow for meaningful pre-vaccination observation periods. **Propensity score stratification or reweighting** can be employed to adjust for differences in baseline demographic, clinical, and healthcare utilization characteristics, as well as temporal factors that may differ across exposure periods. As with trial emulation designs, **negative control outcomes** and **E-values** can be incorporated to assess the potential impact of residual confounding.

As previously suggested, differences between vaccinated and unvaccinated cohorts may be so substantial and systemic that they cannot be fully addressed through conventional confounder adjustment methods. For example, numerous studies have examined whether vaccination increases the risk of all-cause death. While reassuringly none have found any elevated risk, many have reported **implausibly strong protective effects**, with vaccines at times appearing to reduce all-cause mortality by 50% or more, even during the first few weeks post-vaccination before any protective effect against infection-related mortality would have time to take hold (Xu et al., 2023; McCarthy et

al., 2013; McCarthy et al., 2016; Jackson, Nelson, Benson, et al., 2006; Jackson, Jackson, Nelson, et al., 2006; Simonsen et al., 2007; Simonsen et al., 2005).

This phenomenon, known as the **healthy vaccinee effect** (**HVE**), is a form of residual confounding. HVE is particularly pronounced in the immediate post-vaccination period, likely because individuals who are acutely ill or near death are less likely to be vaccinated during those times (McCarthy et al., 2016). As a result, the death rate in recently vaccinated individuals can appear artificially low when compared to historical baselines or unvaccinated subpopulations. The flip side of this phenomenon, evident during the pandemic, is that those who remained unvaccinated or who deferred their next dose were often individuals in poorer health or with limited life expectancy, which could make mortality in that group appear higher than among their age peers. The same selection pattern can also create spuriously low rates of serious outcomes in the weeks leading up to vaccination, biasing that pre-vaccination interval if it is used as a control period.

To mitigate this bias, analyses of post-vaccination mortality should be stratified by time since vaccination, with separate evaluations for the immediate post-vaccination period (e.g., 3-6 weeks) and later time intervals, to help disentangle these effects. Importantly, this form of residual confounding is not limited to mortality analyses or to vaccinated vs. unvaccinated comparisons and may arise in other contexts as well. This underscores the broader need to investigate potential residual confounding through strategies such as the use of negative controls, E-values or other advanced analytic approaches. Note, however, that HVE is only one potential source of residual confounding; it shouldn't be invoked reflexively to dismiss apparent vaccine benefits or to explain away null findings on post-vaccination risk.

Given the limitations described above, unvaccinated controls have been used relatively infrequently in VSD studies. As an alternative, researchers have often turned to **historical controls**, comparing the incidence of specific adverse events in the vaccinated population to background rates observed in prior time periods. A key advantage in this setting is that the historical background rates can be calculated using the same EHR data sources within the VSD medical centers, ensuring a consistent population and with sufficiently granular data to allow for adjustment based on detailed demographic, clinical, and healthcare utilization information.

As previously noted, when using historical controls, it is essential to ensure that the background incidence model is well-calibrated to reflect the demographic and clinical profile of the vaccinated population. However, during periods of rapid change, such as shifts in healthcare utilization, disease dynamics, or population behavior, baseline risks may change in ways that make the historical rates unreliable, even when the model adequately adjusts for key individual-level confounders. For example, for events at higher risk after COVID-19 infection, pre-pandemic historical background rates may be

too low to be accurate during the pandemic, and for events with increasing awareness and diagnostic efforts over time, the historical background rates may be too low to accurately reflect contemporary diagnostic practices. In these situations, the use of **contemporary controls** becomes especially important, as they enable comparisons within the same time frame and help mitigate biases that might arise from relying on outdated background rates.

The limitations of contemporary unvaccinated controls have already been discussed, but **risk-interval methods** provide an alternative strategy that does not require use of unvaccinated or historical controls. These methods only model vaccinated individuals and compare the adverse event incidence between prespecified risk and control periods. The risk period is selected to include the time frame in which the adverse event would be expected to occur if caused by vaccination, while the control period involves other pre-vaccination or post-vaccination time periods. For example, some vaccine safety studies have compared incidence of adverse events during the risk window shortly after vaccination (e.g., day 1 to 42) with a control window at later times (e.g., day 43 and beyond), based on the observation that the vast majority of post-vaccination adverse events occur within the first few days or weeks following immunization. This approach is especially well suited to testing whether a **particular adverse event** is temporally triggered by vaccination.

One widely used risk-interval method is the **Self-Controlled Case Series (SCCS)** design. This approach includes only vaccinated individuals who experienced the event of interest and compares the risk of the event occurring during a defined risk interval relative to that during a control interval. Because each individual serves as their own control, the design inherently adjusts for time-invariant confounders such as individual demographic, clinical, and behavioral characteristics, as well as residual confounders related to the decision to vaccinate or not, such as the healthy vaccinee effect. However, time-varying confounding can still present challenges, particularly when external factors influencing event risk, such as infection rates, seasonal trends, or changes in healthcare delivery, differ substantially between the risk and control periods. As noted earlier, if individuals are less likely to be vaccinated following a specific serious event, the pre-vaccination interval becomes a biased control period for that outcome and therefore should not be used for comparison.

Other risk-interval methods are available in settings for which time-varying confounding is a concern. Klein and colleagues (Klein et al., 2021) used a **concurrent comparator** design in which they compared the adverse event incidence of vaccinated individuals during the risk interval with a matched cohort of vaccinated individuals within the control interval on the same calendar date, including all individuals whether they ever experienced the event or not. Since self-matching was not used, it was still necessary to adjust for potential individual-level confounders between the two groups,

which they accomplished by stratifying the comparison by age, sex, race, VSD site, and calendar date. An alternative **case-centered** (Fireman et al., 2009) analysis only models individuals experiencing the event, comparing the proportion of cases in the risk interval with the expected proportion in the risk interval computed from a matched cohort of vaccinated individuals with the same onset date, age, sex and VSD site (Greene et al., 2012).

A major advantage of these risk-interval methods is that they rely solely on data from vaccinated individuals, making them feasible even in settings with high vaccine coverage. This also helps avoid residual confounding related to vaccination decisions, such as the healthy vaccinee effect or confounding by indication (Salas et al., 1999; Kyriacou & Lewis, 2016) and further avoids any ascertainment bias between vaccinated and unvaccinated individuals. However, interpreting the results of these methods requires caution, as they test whether the incidence of an adverse event **differs between predefined time periods**. This method is effective for detecting early-onset safety signals occurring within days or weeks after vaccination, which are thought to comprise the vast majority of potential vaccine safety issues. However, it may miss potential risks associated with delayed-onset events that fall outside the designated risk window. In such instances, alternative study designs are necessary to assess long-term safety outcomes.

The VSD has implemented a **Rapid Cycle Analysis (RCA)** framework (Lieu et al., 2007) that enables automated, near real-time monitoring of pre-specified adverse events of special interest (AESIs) on a weekly basis. While historical controls have traditionally been used for these analyses, more recent studies have also employed alternative designs, including SCCS and case-centered approaches.

RCA involves repeated statistical testing over time as new data accumulate, which increases the risk of inflated type I error (i.e., false positive findings). As a result, conventional significance thresholds (e.g., $p<0.05$) are not appropriate, and specialized statistical methods are required to control the overall type I error rate in this sequential testing context.

Several approaches have been developed to address this challenge. Wald (1945) first introduced the **sequential probability ratio test** (***SPRT***), which compares the accumulating data against a null and a pre-specified alternative hypothesis for the relative risk (RR). However, its dependence on a fixed, pre-specified targeted relative risk can be limiting. To address this, Kulldorff and collaborators (Kulldorff et al., 2011) proposed the ***MaxSPRT***, an extension that tests a composite alternative hypothesis, eliminating the need to specify a single relative risk. They further developed the ***CmaxSPRT*** (Li & Kulldorff, 2010), which accounts for uncertainty in the expected counts derived from historical baselines. Additionally, Nelson and colleagues (Nelson et al., 2013) introduced group sequential approaches that assume a fixed number of

interim analyses ("looks"), which can offer more statistical power compared to methods allowing unrestricted monitoring.

Other Available Sources of EHR Data for Active Surveillance

Another source of active monitoring for post-approval safety monitoring is the FDA's **Sentinel Initiative**. Its vaccine-focused component, **PRISM** (Post-licensure Rapid Immunization Safety Monitoring program), integrates electronic health records (EHRs) to follow potential safety signals and assess whether they are linked to vaccination. Sentinel also includes the Active Postmarket Risk Identification and Analysis (**ARIA**) system, which harmonizes data from multiple healthcare organizations using a common data model and provides standard analytical programs.

Given the challenges and uncertainties inherent to observational analyses, it is important to validate findings, even if found in well-conducted studies from active monitoring systems like VSD and PRISM. U.S. government sources that can support such validation include datasets from the Veterans Health Administration (**VA**), the Department of Defense (**DOD**), and the Centers for Medicare & Medicaid Services (**CMS**). While these cohorts are not fully representative of the U.S. population, they offer high-value information on specific subgroups, e.g., service members and veterans (VA, DOD), older adults (Medicare), and lower-income populations (Medicaid), and their large, well-curated datasets can accommodate methods similar to those used in VSD studies.

Beyond the United States, national health registries in countries with centralized healthcare systems, such as Canada, the U.K., Israel, and the Nordic countries, and multinational collaborations like the Global Vaccine Data Network (GVDN) (The Global Vaccine Data Network; Faksova et al., 2024) provide additional platforms for vaccine-safety studies and for replicating and validating results from the U.S. active-surveillance systems. Together, these diverse data sources play a critical role in confirming safety signals and assessing their reproducibility across settings.

V-safe: Solicited reporting system through mobile applications

V-safe is a safety monitoring system developed by the CDC during the COVID-19 pandemic to gather real-time post-vaccination data through brief check-ins via a mobile application (V-safe; Myers et al., 2023). Initially designed for COVID-19 vaccines, it has since been expanded to include RSV vaccines and holds potential for broader application in future vaccine monitoring efforts.

Individuals could enroll in V-safe at the time of their initial vaccination. Once enrolled, participants received periodic prompts on their mobile devices to report any post-vaccination symptoms, initially daily during the first week, then weekly and later at longer time intervals. These prompts were intentionally brief, designed to take no more than two minutes to complete, in order to encourage high response rates over time.

Early check-ins focused on common local (e.g., pain or redness at the injection site) and systemic (e.g., fever, fatigue) symptoms. Later prompts inquired about ongoing or new symptoms, whether symptoms interfered with work or daily activities, and whether medical care was sought. Women of childbearing age were also asked to report any positive pregnancy tests.

Unlike VAERS, V-safe has a known number of participants, enabling estimation of incidence rates. Although nonresponse remains a limitation, it is generally less problematic than the reporting bias associated with passive surveillance systems like VAERS. However, generalizability is a concern, as individuals who opt into V-safe, and those who continue to respond, may differ systematically from the broader vaccinated population. Additionally, the limited demographic and clinical information collected by V-safe makes it difficult to adjust for these potential sources of bias or confounding in subsequent analyses.

As with VAERS, V-safe lacks an unvaccinated control group, which limits its ability to determine whether reported events occur more frequently than expected based on background rates. Additionally, the system's design emphasizes immediate, typically mild and transient side effects, rather than serious adverse events (SAEs) or adverse events of special interest (AESIs). Although users can report more serious events in free-text fields, these responses are qualitative and not well suited for systematic analysis.

Improvements to V-safe could include expanding the prompts to explicitly ask about specific SAEs or AESIs. However, doing so would require balancing the need for richer data with the importance of maintaining brevity to preserve engagement and

compliance. Even with such enhancements, V-safe would still face the same fundamental challenge: without internal control groups, it would remain difficult to determine whether the observed event rates exceed what would be expected based on population-level background incidence.

Clinical Immunization Safety Assessment (CISA) project

While VAERS and VSD analyses primarily focus on the epidemiological assessment of potential vaccine risks at the population level, the **Clinical Immunization Safety Assessment (CISA)** project complements these efforts by addressing vaccine safety at the individual clinical level. Established in 2001, CISA is a collaborative network that includes the CDC, eight academic medical centers, and other experts who provide consultation and conduct clinical research to better understand and prevent adverse events following vaccination.

One of CISA's key contributions is its ability to perform clinical case consultations for individual patients, allowing for in-depth evaluations of potential vaccine-related harms. While epidemiologic studies can identify associations between vaccination and adverse events, detailed clinical investigations can uncover underlying biological mechanisms, helping to clarify causality and strengthen the evidence for vaccine-related risk.

Although they focus on individual patients rather than population-level patterns, medical **case reports** serve an important role in the peer-reviewed literature by offering rigorous, structured, and transparent accounts of clinically meaningful or unusual presentations. Many leading medical journals require adherence to **CARE** standards (CARE Case Report Guidelines; Gagnier et al., 2014) standards, which ensure completeness and methodological rigor through clear documentation of the patient's history, clinical findings, diagnostic process, interventions, outcomes, and interpretation.

As noted earlier, in-depth evaluations of individual VAERS reports are sometimes undertaken to reconstruct the circumstances surrounding a reported event and to assess whether a vaccine's contribution was likely, probable, possible, or unlikely. Although the unverified nature of VAERS submissions limits the strength of inferences that can be drawn from them alone, when details are validated and supplemented with

medical records, such investigations can support the development of robust case reports that yield important clinical insights.

Regardless of whether an individual case is identified through a CISA consultation or a VAERS report, it is essential that case report analyses be conducted by clinical researchers with appropriate expertise, and that CARE standards be followed to ensure the evaluation is complete, transparent, and methodologically sound.

The U.S. Vaccine Safety Monitoring System as an Integrated Framework

As emphasized throughout this article, the U.S. vaccine safety monitoring system comprises multiple components, each with distinct strengths and limitations. These components are designed to serve complementary roles and function together as an integrated system to detect and evaluate potential vaccine safety concerns. A comprehensive understanding of vaccine safety emerges only when evidence is considered collectively while integrating all components in the system.

Pre-licensure clinical trials offer rigorous statistical evaluation of common side effects and aggregate measures like the overall rate of serious adverse events. However, due to their limited sample sizes, these trials are typically underpowered for detecting rare but potentially serious adverse events, or downstream events long after vaccination. This limitation is a key reason why the FDA mandates post-approval safety monitoring to identify and characterize any rare risks not detected in the original trials.

VAERS, the Vaccine Adverse Event Reporting System, is a passive surveillance system designed to collect reports of post-vaccination adverse events and identify potential safety signals or hypotheses of harm that warrant further investigation. Its major limitations include reporting bias of unknown magnitude and absence of an appropriate control group, which prevent it from establishing causal relationships or estimating incidence rates. Nevertheless, with careful analysis that accounts for these limitations, VAERS can be an effective early warning system for identifying potential safety signals.

V-safe is a newer system that solicits post-vaccination symptom reports from enrolled participants through periodic mobile prompts. Because the number of participants is known, V-safe allows estimation of incidence rates. However, it primarily captures mild and transient side effects, along with general indicators such as whether symptoms

interfered with daily activities (i.e., were severe) or required medical care (i.e., were serious). It does not directly collect detailed data on specific serious adverse events (SAE) or adverse events of special interest (AESI) and it lacks a control group, limiting its ability to assess whether reported events occur more frequently than expected.

The CDC's Vaccine Safety Datalink (**VSD**) and FDA's Post-licensure Rapid Immunization Safety Monitoring (**PRISM**) within the Sentinel system are active surveillance systems that use electronic health records to monitor medically verified adverse events. Unlike passive surveillance systems, they avoid reporting bias and enable scientifically rigorous comparisons using various types of control groups. This allows for more definitive assessment of whether specific adverse events occur at elevated rates following vaccination yet still are prone to confounding and ascertainment bias so must be rigorously analyzed using state-of-the-art epidemiological methods. These systems are used both to investigate safety signals identified in VAERS and to conduct weekly Rapid Cycle Analysis (RCA) that monitors pre-specified AESIs in near real time. While they only capture adverse events that prompt medical attention, this makes them particularly well-suited for monitoring the serious events of greatest concern.

Safety signals identified in VSD or PRISM can be further evaluated for consistency using other large-scale electronic health record systems, such as those maintained by the Veterans Health Administration (VA), Department of Defense (DOD), and Centers for Medicare & Medicaid Services (CMS), or through international data sources such as national registries in countries with centralized healthcare systems and collaborative networks like the Global Vaccine Data Network (GVDN). Reproducibility of findings across multiple systems strengthens confidence in identified safety signals.

Finally, the Clinical Immunization Safety Assessment (**CISA**) project provides expert clinical consultation and in-depth case investigations for individual patients. These investigations can help validate epidemiological findings, offer clinical confirmation of vaccine-related harm, and explore possible biological mechanisms.

By integrating evidence from all of these complementary components of the system, we gain a more complete and accurate understanding of potential vaccine safety issues. Each component plays a vital role, and their combined insights contribute to create a more robust and reliable safety monitoring infrastructure.

Suggested Improvements to the U.S. Safety Monitoring System

While the various components of the U.S. safety monitoring system together form a broad and robust infrastructure, each has its own limitations, and the overall system can be substantially strengthened. Below, I outline several suggested improvements to enhance the effectiveness of individual components and the system as a whole, some of which were discussed in more detail in the respective sections.

Pre-licensure Clinical Trials

- **Increase size and duration.** Larger, longer trials improve power to detect uncommon serious adverse events, severe infections, and deaths, even though very rare events may still escape detection.
- **Use more representative cohorts.** Include wider age ranges and comorbidity profiles to more closely mirror the populations that will ultimately receive the vaccine.
- **Expand prespecified AESIs.** Define a broader list of adverse events of special interest and, where appropriate, add more sensitive assays (e.g., troponin assays for cardiac injury).
- **Consider delayed-vaccination designs.** When placebo control is considered impractical and disease exposure not imminent, randomizing participants to immediate vs. delayed schedules allows causal inference during the delay period without denying any participants the vaccine.

VAERS (Passive Surveillance)

- **Boost reporting.** Explore strategies to increase reporting, including piloting automatic EHR-to-VAERS submission within healthcare systems to capture more events.
- **Quantify reporting bias.** Conduct studies to estimate underreporting factors (URF) for key events and track how bias varies by vaccine, event, timing, and context.
- **Develop granular baseline risk models.** Build dynamic models for key adverse events stratified by demographic and clinical factors and estimating trends over calendar time to generate more accurate contemporary background rates.

- **Incorporate underreporting factors (URF)**. Either (i) compute the minimum URF that would trigger a safety signal or (ii) embed URF uncertainty in a Bayesian framework and test sensitivity to its specification.
- **Improve public education.** Provide clear guidance so the system is used appropriately and misinterpretation is minimized.

V-safe (Solicited Reporting)

- **Broaden enrollment.** Continue to expand its use, especially for any adult vaccines and boosters.
- **Collect richer covariates.** Add basic demographic and clinical information to support more rigorous adjustment for potential bias.
- **Solicit key AESIs.** Include a concise set of specific serious outcomes in prompts to capture high-priority serious events more systematically.

Active surveillance systems (VSD, PRISM, VA, DOD, CMS)

- **Utilize more cutting-edge designs.** Use more innovative observational study designs (e.g., trial emulation, time-varying exposure models).
- **Strengthen confounder control.** Improve adjustment methods and systematically evaluate residual confounding.
- **Use multiple analytic lenses.** Apply and cross-check findings across alternative design and analytical approaches to assess the robustness of the method used.
- **Validate across datasets.** Replicate signals seen in VSD or PRISM within VA, DOD, and CMS data, and collaborate with international networks (e.g., GVDN) and countries with centralized medical systems for external validation.
- **Better address long-term risks.** Develop study designs tailored to identifying potential longer-term risks, which current, commonly used designs are not well-suited to detect.
- **Expand RCA.** Continue weekly Rapid Cycle Analyses (RCA) of Adverse Events of Special Interest (AESI) and expand to include additional AESIs and SAEs.
- **Better coordinate CDC and FDA efforts:** Coordinate the CDC's VSD and FDA's PRISM so that they function as an integrated system, with shared protocols, harmonized case definitions, joint signal triage, and reciprocal replication, rather than operating in separate silos.
- **Improve public awareness.** Better publicize these active-surveillance systems, the studies they produce, and the RCA system, communicating results

transparently to build public understanding and confidence that vaccine safety is being rigorously assessed.

Clinical validation (CISA)

- **More systematic follow-up.** Conduct regular, structured clinical follow-up for validated safety signals to deepen clinical characterization and help identify potential mechanisms, and to clinically validate and deeply characterize individual reports of serious events in VAERs, including deaths.
- **More explicit assessment of case reports and common or high-profile anecdotes.** Establish a process to monitor published clinical case reports that suggest potential safety signals, especially for events producing numerous case reports, and evaluate these events within the passive and active monitoring systems. Also track anecdotal reports from media and social media; when patterns emerge, assess them across the multicomponent safety monitoring system to detect unknown risks and reassure the public that the system is responsive and concerns are being investigated and addressed.

Transparent Communication

- **Better educate the public about the safety system:** Clearly describe and summarize the multicomponent safety monitoring framework to build understanding and confidence.
- **Better disclose safety-monitoring efforts:** Communicate transparently what is being done: report hypothesis-generation from case reports and passive systems, validation in active systems, and any subsequent attempts to reproduce findings in other datasets.

Conclusion

Vaccines have eradicated or contained once-common childhood diseases, and adult vaccines protect vulnerable adults from severe outcomes. As memories of these past harms fade, childhood vaccination can seem like a **Chesterton's Fence** (Parrish 2020), a precaution whose value is easy to overlook, which if changed without sufficient deliberation could invite unintended negative consequences. This fading memory has also heightened attention to vaccine risks and potential harms.

Identifying post-vaccine risks is inherently challenging, particularly when the risks are rare but potentially serious. Randomized clinical trials provide rigorous estimates for common side effects and aggregate measures such as overall serious adverse events, but their size and duration limit power for detecting rare events and make them poorly suited to uncover potential long-term effects years after vaccination. These limitations are a principal reason why the FDA mandates post-approval safety monitoring to detect and characterize risks not seen in pre-licensure trials.

The U.S. vaccine safety monitoring system comprises complementary components, each with its own distinct strengths and limitations, designed to function in coordination. Clinical case reports and passive surveillance (e.g., VAERS) generate hypotheses by flagging potential safety signals. These signals are then tested and validated in active reporting systems (e.g., the CDC's VSD and FDA's PRISM), and ideally replicated in other large U.S. data sets (e.g., VA, DOD, CMS), integrated health systems abroad, and international networks (e.g., GVDN), with clinical networks (e.g., CISA) providing case-level review, mechanistic insight, and guidance for clinicians.

Although the COVID-19 pandemic drew greater public attention to vaccine safety, and to parts of the monitoring system, there is still limited understanding of the system's full architecture, how each component operates, and how they are intended to interact to form an integrated safety net. This white paper has reviewed the major components, their roles, strengths, limitations, and analytical considerations, and explained how they are designed to function together within a comprehensive monitoring framework. It also proposes improvements, both component-specific and system-wide, to strengthen detection and evaluation of vaccine-related risks. The aim is to improve public explanations of the system, spur discussion on how to enhance it, and help sustain a robust safety monitoring infrastructure that earns public confidence.

References

Almadani, O. A., & Alshammari, T. M. (2022). Vaccine adverse event reporting system (VAERS): Evaluation of 31 years of reports and pandemics' impact. *Saudi Pharmaceutical Journal*, *30*(12), 1725–1735. https://doi.org/10.1016/j.jsps.2022.10.001

Austin, P. C. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, *46*(3), 399–424. https://doi.org/10.1080/00273171.2011.568786

Banks, D., Woo, E. J., Burwen, D. R., Perucci, P., Braun, M. M., & Ball, R. (2005). Comparing data mining methods on the VAERS database. *Pharmacoepidemiology and Drug Safety*, *14*(9), 601–609. https://doi.org/10.1002/pds.1107

Barosa, M., & Prasad, V. (2025). Characteristics of vaccine safety observational studies and authors' attitudes: A systematic review. *The American Journal of Medicine*, *138*(2), 254–261.e5. https://doi.org/10.1016/j.amjmed.2024.10.007

Braun, M. M., & Ellenberg, S. S. (1997). Descriptive epidemiology of adverse events after immunization: reports to the Vaccine Adverse Event Reporting System (VAERS), 1991-1994. *The Journal of Pediatrics*, *131*(4), 529–535. https://doi.org/10.1016/s0022-3476(97)70056-8

CARE Case Report Guidelines. *What are the CARE guidelines?* https://www.care-statement.org/

Centers for Disease Control and Prevention. (2025). About the Vaccine Safety Datalink (VSD). *Vaccine Safety Systems*. https://www.cdc.gov/vaccine-safety-systems/vsd/index.html

Chen, R. T., Glasser, J. W., Rhodes, P. H., Davis, R. L., Barlow, W. E., Thompson, R. S., Mullooly, J. P., Black, S. B., Shinefield, H. R., Vadheim, C. M., Marcy, S. M., Ward, J. I., Wise, R. P., Wassilak, S. G., & Hadler, S. C. (1997). Vaccine Safety Datalink project: a new tool for improving vaccine safety monitoring in the United States. *Pediatrics*, *99*(6), 765–773. https://doi.org/10.1542/peds.99.6.765

DuMouchel, W. (1999). Bayesian data mining in large frequency tables, with an application to the FDA Spontaneous Reporting System. *The American Statistician*, *53*(3), 177–190. https://doi.org/10.1080/00031305.1999.10474456

DuMouchel, W., & Pregibon, D. (2001). Empirical Bayes screening for multi-item associations. *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 67–76. https://doi.org/10.1145/502512.502526

Eberth, J. M., Kline, K. N., Moskowitz, D. A., Montealegre, J. R., & Scheurer, M. E. (2014). The role of media and the Internet on vaccine adverse event reporting: a case study of human papillomavirus vaccination. *The Journal of Adolescent Health*, *54*(3), 289–295. https://doi.org/10.1016/j.jadohealth.2013.09.005

Evans, S. J., Waller, P. C., & Davis, S. (2001). Use of proportional reporting ratios (PRRs) for signal generation from spontaneous adverse drug reaction reports. *Pharmacoepidemiology and Drug Safety*, *10*(6), 483–486. https://doi.org/10.1002/pds.677

Faksova, K., Walsh, D., Jiang, Y., Griffin, J., Phillips, A., Gentile, A., Kwong, J. C., Macartney, K., Naus, M., Grange, Z., Escolano, S., Sepulveda, G., Shetty, A., Pillsbury, A., Sullivan, C., Naveed, Z., Janjua, N. Z., Giglio, N., Perälä, … Hviid, A. (2024). COVID-19 vaccines and adverse events of special interest: A multinational Global Vaccine Data Network (GVDN) cohort study of 99 million vaccinated individuals. *Vaccine*, *42*(9), 2200–2211. https://doi.org/10.1016/j.vaccine.2024.01.100

Finney, D. J. (1974). Systemic signalling of adverse reactions to drugs. *Methods of Information in Medicine*, *13*(1), 1–10.

Fireman, B., Lee, J., Lewis, N., Bembom, O., van der Laan, M., & Baxter, R. (2009). Influenza vaccination and mortality: differentiating vaccine effects from bias. *American Journal of Epidemiology*, *170*(5), 650–656. https://doi.org/10.1093/aje/kwp173

Gagnier, J. J., Kienle, G., Altman, D. G., Moher, D., Sox, H., Riley, D., & CARE Group. (2013). The CARE guidelines: Consensus-based clinical case reporting guideline Development. *Global Advances in Health and Medicine*, *2*(5), 38–43. https://doi.org/10.7453/gahmj.2013.008

Gee, J., Shimabukuro, T. T., Su, J. R., Shay, D., Ryan, M., Basavaraju, S. V., Broder, K. R., Clark, M., Creech, C.B., Cunningham, F., Goddard, K., Guy, H., Edwards, K. M., Forshee, R., Hamburger, T., Hause, A. M., Klein, N. P., Kracalik, I., Lamer, C., … & Anderson, S. (2024). Overview of U.S. COVID-19 vaccine safety surveillance systems. *Vaccine*, *42*(Suppl 3), 125748. https://doi.org/10.1016/j.vaccine.2024.02.065

Greene, S. K., Rett, M., Weintraub, E. S., Li, L., Yin, R., Amato, A. A., Ho, D. T., Sheikh, S. I., Fireman, B. H., Daley, M. F., Belongia, E. A., Jacobsen, S. J., Baxter, R., Lieu, T. A., Kulldorff, M., Vellozzi, C., & Lee, G. M. (2012). Risk of confirmed Guillain-Barre syndrome following receipt of monovalent inactivated influenza A (H1N1) and seasonal influenza vaccines in the Vaccine Safety Datalink Project, 2009-2010. *American Journal of Epidemiology*, *175*(11), 1100–1109. https://doi.org/10.1093/aje/kws195

Google. *Google Trends*. https://trends.google.com/trends/explore?date=2019-01-01%202022-12-31&geo=US&q=vaers,Vaccine%20safety%20datalink&hl=en

Hernán, M. A., Wang, W., & Leaf, D. E. (2022). Target trial emulation: A framework for causal inference from observational data. *JAMA*, *328*(24), 2446–2447. https://doi.org/10.1001/jama.2022.21383

Hill, A.B. (1965). The environment and disease: Association or causation? *Proceedings of the Royal Society of Medicine*, *58*(5), 295-300. https://doi.org/10.1177/003591576505800503

Jackson, L. A., Jackson, M. L., Nelson, J. C., Neuzil, K. M., & Weiss, N. S. (2006). Evidence of bias in estimates of influenza vaccine effectiveness in seniors. *International Journal of Epidemiology*, *35*(2), 337–344. https://doi.org/10.1093/ije/dyi274

Jackson, L. A., Nelson, J. C., Benson, P., Neuzil, K. M., Reid, R. J., Psaty, B. M., Heckbert, S. R., Larson, E. B., & Weiss, N. S. (2006). Functional status is a confounder of the association of influenza vaccine and risk of all cause mortality in seniors. *International Journal of Epidemiology*, *35*(2), 345–352. https://doi.org/10.1093/ije/dyi275

Klein, N. P., Lewis, N., Goddard, K., Fireman, B., Zerbo, O., Hanson, K. E., Donahue, J. G., Kharbanda, E. O., Naleway, A., Nelson, J. C., Xu, S., Yih, W. K., Glanz, J. M., Williams, J. T. B., Hambidge, S. J., Lewin, B. J., Shimabukuro, T. T., DeStefano, F., & Weintraub,

E. S. (2021). Surveillance for adverse events after COVID-19 mRNA vaccination. *JAMA*, *326*(14), 1390–1399. https://doi.org/10.1001/jama.2021.15072

Kulldorff, M., Davis, R. L., Kolczak†, M., Lewis, E., Lieu, T., & Platt, R. (2011). A maximized sequential probability ratio test for drug and vaccine safety surveillance. *Sequential Analysis*, *30*(1), 58–78. https://doi.org/10.1080/07474946.2011.539924

Kyriacou, D. N., & Lewis, R. J. (2016). Confounding by indication in clinical research. *JAMA*, *316*(17), 1818–1819. https://doi.org/10.1001/jama.2016.16435

Lazarus, R. (2010). Electronic Support for Public Health - Vaccine Adverse Event Reporting System (ESP:VAERS) - Final Report. (Prepared by Harvard Pilgrim Health Care, Inc. under Grant No. R18 HS017045). Rockville, MD: Agency for Healthcare Research and Quality. https://digital.ahrq.gov/ahrq-funded-projects/electronic-support-public-health-vaccine-adverse-event-reporting-system/final-report

Li, L., & Kulldorff, M. (2010). A conditional maximized sequential probability ratio test for pharmacovigilance. *Statistics in Medicine*, *29*(2), 284–295. https://doi.org/10.1002/sim.3780

Lieu, T. A., Kulldorff, M., Davis, R. L., Lewis, E. M., Weintraub, E., Yih, K., Yin, R., Brown, J. S., Platt, R., & Vaccine Safety Datalink Rapid Cycle Analysis Team (2007). Real-time vaccine safety surveillance for the early detection of adverse events. *Medical Care*, *45*(10 Supl 2), S89–S95. https://doi.org/10.1097/MLR.0b013e3180616c0a

Lipsitch, M., Tchetgen Tchetgen, E., & Cohen, T. (2010). Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology*, 21(3), 383–388. https://doi.org/10.1097/EDE.0b013e3181d61eeb

Miller, E. R., McNeil, M. M., Moro, P. L., Duffy, J., & Su, J. R. (2020). The reporting sensitivity of the Vaccine Adverse Event Reporting System (VAERS) for anaphylaxis and for Guillain-Barré syndrome. *Vaccine*, *38*(47), 7458–7463. https://doi.org/10.1016/j.vaccine.2020.09.072

McCarthy, N. L., Gee, J., Sukumaran, L., Weintraub, E., Duffy, J., Kharbanda, E. O., Baxter, R., Irving, S., King, J., Daley, M. F., Hechter, R., & McNeil, M. M. (2016). Vaccination and 30-day mortality risk in children, adolescents, and young adults. *Pediatrics*, *137*(3), e20152970. https://doi.org/10.1542/peds.2015-2970

McCarthy, N. L., Weintraub, E., Vellozzi, C., Duffy, J., Gee, J., Donahue, J. G., Jackson, M. L., Lee, G. M., Glanz, J., Baxter, R., Lugg, M. M., Naleway, A., Omer, S. B., Nakasato, C., Vazquez-Benitez, G., & DeStefano, F. (2013). Mortality rates and cause-of-death patterns in a vaccinated population. *American Journal of Preventive Medicine*, *45*(1), 91–97. https://doi.org/10.1016/j.amepre.2013.02.020

McNeil, M. M., Gee, J., Weintraub, E. S., Belongia, E. A., Lee, G. M., Glanz, J. M., Nordin, J. D., Klein, N. P., Baxter, R., Naleway, A. L., Jackson, L. A., Omer, S. B., Jacobsen, S. J., & DeStefano, F. (2014). The Vaccine Safety Datalink: successes and challenges monitoring vaccine safety. *Vaccine*, *32*(42), 5390–5398. https://doi.org/10.1016/j.vaccine.2014.07.073

McNeil, M. M., Li, R., Pickering, S., Real, T. M., Smith, P. J., & Pemberton, M. R. (2013). Who is unlikely to report adverse events after vaccinations to the Vaccine Adverse Event Reporting System (VAERS)? *Vaccine*, *31*(24), 2673–2679. https://doi.org/10.1016/j.vaccine.2013.04.009

Meyer, S. (2025). Update on CDC's Covid-19 vaccine safety monitoring. https://www.cdc.gov/acip/downloads/slides-2025-06-25-26/04-Meyer-COVID-508.pdf. Accessed on December 3, 2025.

Myers, T. R., Marquez, P. L., Gee, J. M., Hause, A. M., Panagiotakopoulos, L., Zhang, B., McCullum, I., Licata, C., Olson, C. K., Rahman, S., Kennedy, S. B., Cardozo, M., Patel, C. R., Maxwell, L., Kallman, J. R., Shay, D. K., & Shimabukuro, T. T. (2023). The v-safe after vaccination health checker: Active vaccine safety monitoring during CDC's COVID-19 pandemic response. *Vaccine*, *41*(7), 1310–1318. https://doi.org/10.1016/j.vaccine.2022.12.031

Nelson, J. C., Yu, O., Dominguez-Islas, C. P., Cook, A. J., Peterson, D., Greene, S. K., Yih, W. K., Daley, M. F., Jacobsen, S. J., Klein, N. P., Weintraub, E. S., Broder, K. R., & Jackson, L. A. (2013). Adapting group sequential methods to observational postlicensure vaccine safety surveillance: results of a pentavalent combination DTaP-IPV-Hib vaccine safety study. *American journal of epidemiology*, *177*(2), 131–141. https://doi.org/10.1093/aje/kws317

Parrish, S. (2020, March 10). Chesterton's fence: A lesson in thinking. *Farnam Street*. https://fs.blog/chestertons-fence/

Prasad, V. (2025). Deaths in children due to COVID-19 vaccines and CBER's path forward. Leaked internal email posted by the Washington Post.

https://www.washingtonpost.com/documents/3042d15c-676b-48ac-8148-1a2204ef420e.pdf

Rosenthal, S., & Chen, R. (1995). The reporting sensitivities of two passive surveillance systems for vaccine adverse events. *American Journal of Public Health*, *85*(12), 1706–1709. https://doi.org/10.2105/ajph.85.12.1706

Salas, M., Hofman, A., & Stricker, B. H. (1999). Confounding by indication: an example of variation in the use of epidemiologic terminology. *American Journal of Epidemiology*, *149*(11), 981–983. https://doi.org/10.1093/oxfordjournals.aje.a009758

Sarntivijai, S., Xiang, Z., Shedden, K. A., Markel, H., Omenn, G. S., Athey, B. D., & He, Y. (2012). Ontology-based combinatorial comparative analysis of adverse events associated with killed and live influenza vaccines. *PLOS One*, *7*(11), e49941. https://doi.org/10.1371/journal.pone.0049941

Scholl, J., van Hunsel, F., & van Puijenbroek, E. (2025). Timing matters: Exploring the role of the time to onset in recall bias for Adverse Events Following Immunization (AEFIs) of COVID-19 vaccines from spontaneous reports. *Pharmacoepidemiology*, *4*(2), 8. https://doi.org/10.3390/pharma4020008

Shimabukuro, T. T., Nguyen, M., Martin, D., & DeStefano, F. (2015). Safety monitoring in the Vaccine Adverse Event Reporting System (VAERS). *Vaccine*, *33*(36), 4398–4405. https://doi.org/10.1016/j.vaccine.2015.07.035

Simonsen, L., Reichert, T. A., Viboud, C., Blackwelder, W. C., Taylor, R. J., & Miller, M. A. (2005). Impact of influenza vaccination on seasonal mortality in the US elderly population. *Archives of Internal Medicine*, *165*(3), 265–272. https://doi.org/10.1001/archinte.165.3.265

Simonsen, L., Taylor, R. J., Viboud, C., Miller, M. A., & Jackson, L. A. (2007). Mortality benefits of influenza vaccination in elderly people: an ongoing controversy. *The Lancet Infectious Diseases*, *7*(10), 658–666. https://doi.org/10.1016/S1473-3099(07)70236-0

Sukumaran, L., McCarthy, N. L., Li, R., Weintraub, E. S., Jacobsen, S. J., Hambidge, S. J., Jackson, L. A., Naleway, A. L., Chan, B., Tao, B., & Gee, J. (2015). Demographic characteristics of members of the Vaccine Safety Datalink (VSD): A comparison with the United States population. *Vaccine*, *33*(36), 4446–4450. https://doi.org/10.1016/j.vaccine.2015.07.037

Szabo, L. (2025, December 1; updated 2025, December 2). FDA official proposes 'impossible' standards for vaccine testing that could curtail access to immunizations. The Center for Infectious Disease Research and Policy (CIDRAP), University of Minnesota. https://www.cidrap.umn.edu/childhood-vaccines/fda-official-proposes-impossible-standards-vaccine-testing-could-curtail-access. Accessed on December 3, 2025.

The Global Vaccine Data Network. *News*. https://www.globalvaccinedatanetwork.org/news

The Vaccine Adverse Event Reporting System (VAERS) Request. https://wonder.cdc.gov/vaers.html

VAERS. *Report an Adverse Event to VAERS*. https://vaers.hhs.gov/reportevent.html

VanderWeele, T.J., & Ding, P. (2017). Sensitivity analysis in observational research: Introducing the E-Value. *Annals of Internal Medicine*,*167*(4), 268-274. https://doi.org/10.7326/M16-2607

Verstraeten, T., Baughman, A. L., Cadwell, B., Zanardi, L., Haber, P., Chen, R. T., & Vaccine Adverse Event Reporting System Team. (2001). Enhancing vaccine safety surveillance: a capture-recapture analysis of intussusception after rotavirus vaccination. *American Journal of Epidemiology*, *154*(11), 1006–1012. https://doi.org/10.1093/aje/154.11.1006

V-safe. https://vsafe.cdc.gov/home

Wald, A. (1945). Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, *16*(2), 117–186. http://www.jstor.org/stable/2235829

Xu, S., Huang, R., Sy, L. S., Hong, V., Glenn, S. C., Ryan, D. S., Morrissette, K., Vazquez-Benitez, G., Glanz, J. M., Klein, N. P., Fireman, B., McClure, D., Liles, E. G., Weintraub, E. S., Tseng, H. F., & Qian, L. (2023). A safety study evaluating non-COVID-19 mortality risk following COVID-19 vaccination. *Vaccine*, *41*(3), 844–854. https://doi.org/10.1016/j.vaccine.2022.12.036