From Criteria to Consequence: Toward a Pre-Specified Escalation Rule in Post-Market Safety Governance

George Kujo (九条丈二)

Independent Researcher, Tokyo, Japan

ORCID: https://orcid.org/0009-0006-6864-0145

george.kujo.irjp@gmail.com

Working paper — not peer reviewed

Version 6 · May 2026

Abstract

Post-market safety governance pre-specifies a great deal. Active surveillance monitors pre-specified outcomes against statistical thresholds designed for repeated testing. European signal-management rules fix deadlines in days, do not require causality before obligation attaches, require that rationale be recorded, and publish the resulting recommendations — including recommendations that no further action is needed. The considerations bearing on those recommendations are themselves set out in advance: severity, reversibility, strength of evidence, exposed population, vulnerable groups, the consequences of discontinuation.

What is not set out in advance, in the instruments examined here, is the rule connecting those considerations to an outcome. Which combination of them should issue in which response is settled in the course of assessing each signal. Reasons are given, and published, but there is no term stated beforehand against which a reader can ask whether the conclusion departed from what had been undertaken.

This paper asks whether that rule should be published in advance, and proposes a specific form: not a trigger compelling a measure whenever a condition is met, but a rebuttable presumption. A published rule raises a presumption of a defined response; a regulator may depart from it, provided the departure is published, expressed against the rule, and — where the ground of departure is one the rule did not anticipate — added to the rule prospectively. What such an arrangement removes is not discretion but the freedom to settle the standard after the event it was to govern.

Nuclear operating limits and drinking water notification rules show that binding a stated condition to a stated consequence is institutionally achievable in complex domains, though not that a rule of this kind could be calibrated for pharmacovigilance. Two cases show how escalation proceeds without one. The paper sets out three candidate forms such a rule might take, together with what a departure would have to contain, and leaves substantial questions open — including the decision rule under statistical uncertainty and the risk that pre-commitment relocates discretion into measurement. It argues that the arrangement merits evaluation, not that it should be adopted.

Keywords: post-market surveillance; regulatory escalation; rebuttable presumption; pharmacovigilance; auditability; safety governance

1. Introduction

1.1 The question

Any safety-critical system operated at scale will eventually generate a signal that was not anticipated at approval. Detecting such signals is a mature activity: linked databases, continuous monitoring and sequential statistical methods are in routine use. Detection is not thereby reliable — ascertainment, denominators, unstable background rates and subgroup discovery all remain hard, and Section 6 returns to them — but the machinery exists and is operated.

The question here concerns what happens after. Once a signal has completed whatever evaluation a jurisdiction requires, what determines the response, and is that determinant stated before the signal or after it?

A note on terminology is needed, since one term is used differently in different places. In European pharmacovigilance a confirmed signal is one a rapporteur has decided requires further analysis; it does not mean causality has been established. In active surveillance a signal is a test statistic exceeding a pre-specified threshold. Neither implies a settled causal question. This paper uses assessed signal for a signal that has completed formal evaluation and is ready for a regulatory recommendation, reserving jurisdiction-specific terms for the passages describing those jurisdictions.

1.2 Scope, method and status

The argument is offered in support of, not in opposition to, the institutions it examines. The illustration concerns COVID-19 vaccination; nothing here argues against vaccination or claims that any particular decision was wrong.

The method is documentary comparison of regulatory instruments, for conceptual rather than causal purposes. Four domains were selected purposively: two known to bind post-deployment conditions to consequences, two in which the step from assessment to response appeared discretionary. Such a design can show that an arrangement is feasible somewhere and that domains vary in adopting it. It cannot identify what causes the variation, and no claim of that kind is made. Nor can the nuclear and drinking water comparators show that a rule of the kind proposed in Section 6 could be calibrated for pharmacovigilance; they establish only the general institutional possibility of binding a stated condition to a stated consequence. No systematic search of any jurisdiction's regulatory corpus has been undertaken. Every statement that something was not found is confined to the instruments and meeting materials cited, and should be read as defeasible.

This is a working paper deposited as a public record; it has not been peer reviewed. Earlier versions were submitted to several journals in April and May 2026 and were not accepted, and successive rounds of review produced substantial changes. Claims that have been withdrawn include an asserted asymmetry between the harms and benefits of preventive intervention, specific numerical ceilings that had no derivation, and a characterisation of post-market pharmaceutical regulation as lacking pre-specification. Citation errors were corrected and an earlier preprint deposit withdrawn.

2. What is already pre-specified

Binding a numerical condition to a consequence in advance is sometimes treated as unworkable where the system is complex and the evidence incomplete. Two regimes show otherwise. A third body of practice, in pharmacovigilance itself, shows how much is already pre-specified there — which is why the gap must be located with care.

2.1 Nuclear: limiting conditions for operation

United States nuclear regulation requires each operating licence to contain technical specifications, and those to contain limiting conditions for operation: the lowest functional capability required for safe operation. Where a limiting condition is not met, 10 CFR 50.36(c)

(2) provides that the licensee shall shut down the reactor or follow such remedial action as the specifications permit until the condition can be met.

Three elements compose the structure: a threshold fixed at licensing; a completion time running from the deviation, within which the condition must be restored; and a mandatory consequence if that time expires, with the schedule for each shutdown mode specified in hours. Two features matter here. The consequence is not weighed at the moment of deviation; that weighing was done when the specifications were approved. And guidance developed for these specifications provides, in the contexts to which it applies, that formal cause evaluation is not required before the required action is taken, because time is short. Diagnosis and response are separated, and response does not wait.

2.2 Drinking water: tiered notification on fixed deadlines

The Public Notification Rule, at 40 CFR Part 141 Subpart Q, matches urgency to seriousness across three tiers. Tier 1, for violations with potential to cause serious adverse health effects from short-term exposure, requires public notice within twenty-four hours and consultation with the primacy agency in the same window; Tier 2 carries thirty days, Tier 3 one year.

The trigger is exceedance of a maximum contaminant level fixed in advance. The deadline is stated in hours rather than qualified by reasonableness. The content of the notice is prescribed in ten elements — the nature of the violation, potential health effects, the population at particular risk — so that the obligation cannot be discharged by an announcement conveying nothing. And where a required consultation does not occur within its window, certain violations are elevated automatically: the default on procedural failure is escalation rather than lapse.

2.3 Pharmacovigilance and active surveillance

Post-market pharmaceutical surveillance is not a domain in which little is fixed in advance. Four kinds of pre-specification are already in place.

First, detection thresholds. Active surveillance systems monitor pre-specified outcomes in pre-specified post-vaccination risk intervals, analysed at weekly or biweekly cycles against an expected or comparison rate. Because repeated testing would otherwise generate false positives, sequential methods — the sequential probability ratio test and its maximised variant — were developed and are applied, so that a statistic exceeding a threshold constitutes

a signal. Outcomes, intervals, threshold and method are all fixed before the data are examined. What the threshold triggers, however, is investigation rather than any regulatory measure.

Second, procedural deadlines. Under the European Union’s good pharmacovigilance practices, Module IX governs signal management as a staged process with several stages timed. An emerging safety issue must be notified within three working days. A standalone signal notification is to be made as soon as possible and no later than thirty days after the marketing authorisation holder concludes that further analysis is needed; signal confirmation is to be taken within thirty days of receipt. Confirmation does not require causality to be established.

Third, reason-giving, both internal and public. Signal-management activities including analyses, decisions and their rationale must be recorded in an audit trail, and where a validated signal is not confirmed the justification must be communicated to the Agency and the Committee. Beyond that, the recommendations issued following assessment are published, and they include conclusions that no further evaluation or action is needed at present. Reasons are neither withheld nor unannounced.

Fourth — and this is the element most easily overlooked — the considerations bearing on the decision are themselves specified in advance. The same guidance directs that prioritisation take account of the severity, seriousness and outcome of the reaction, its reversibility, the strength of the evidence, the size of the exposed population, the presence of vulnerable groups, the consequences of discontinuing treatment, and the regulatory measures that might be contemplated. These are not composed after the signal appears. They are published, and they are broadly the considerations any competent analyst would identify.

2.4 What is missing

The nuclear and drinking water regimes share a property that the elements listed above, taken separately, do not supply. In both, a defined condition is bound in advance to a defined consequence within a defined period. Crossing the limit does not merely open an assessment whose outcome is unconstrained; it determines, beforehand, what is then to happen.

It follows that the gap in post-market pharmaceutical governance is not an absence of criteria. The criteria are published, and they are the right criteria. What was not identified in the

instruments examined here is the rule connecting them: given an assessed signal with a particular severity, reversibility, evidentiary strength, absolute and relative excess, and exposed population, no published statement indicates which response is expected, or what would count as a departure from expectation. The considerations are fixed in advance; their aggregation into an outcome is not. Module IX is explicit on the point in describing how statistical output is to be treated: not every statistical alert requires investigation, some non-statistical findings do, and scientific and clinical judgment is paramount.

This distinction also governs what can be said about reason-giving. Recommendations are published, and a recommendation that no further action is needed is itself a public statement. But a published conclusion is not a justification measured against a standard that existed before the signal. Where the mapping from considerations to outcome is settled while the signal is being assessed, the reasons given are coherent by construction, and a reader has no independent term against which to test them. That, rather than secrecy, is the deficiency identified here.

3. Two cases

3.1 The Boeing 737 MAX

Aviation is routinely offered as the exemplar of engineering safety governance, and at certification the description is apt. The 737 MAX shows that the strength did not extend to the post-deployment decision.

Lion Air Flight 610 crashed on 29 October 2018, killing 189. Ethiopian Airlines Flight 302 crashed on 10 March 2019, killing 157. Between them, four months and ten days, the type flew commercially worldwide. After the second accident, China ordered suspension on 11 March; several jurisdictions including the European Union Aviation Safety Agency acted on 12 March; the Federal Aviation Administration issued its order on 13 March, last among major regulators, having initially affirmed the type’s continued airworthiness on the ground that the evidence was insufficient. The grounding lasted 619 days, and the type returned to service after design changes.

Nothing in the certification basis stated that a given number of hull losses in given circumstances within a given period would raise a presumption of suspension. Had such a rule existed, the question after the second accident would have been whether its conditions

were met, and if so whether departure was justified — narrow and answerable questions — rather than whether the aircraft should be grounded, which is broad and contestable. Authorities examining the same two accidents reached different conclusions on different days. It should not be assumed that the absence of a rule explains the divergence; differences in the evidence reaching each authority and in institutional risk tolerance are equally available, and this paper cannot distinguish among them. What the episode does establish is that each authority decided through its own judgment rather than by applying a common published standard, and that neither the timing nor the divergence could be assessed against one.

3.2 The 2021 myocarditis signal in Japan

Japan’s handling of the 2021 mRNA vaccine myocarditis signal is examined because the authorities did act, and differentiated by age and sex. That sharpens the question: not whether a regulator responded, but on what stated basis the timing and successive stages were determined.

Post-market vaccine safety in Japan is reviewed by a subcommittee on adverse reactions of the Health Sciences Council, meeting jointly with a safety measures investigation committee of the Pharmaceutical Affairs and Food Sanitation Council. Myocarditis-related material appears in the published papers of successive meetings through 2021, at intervals of roughly two to three weeks. Three of those meetings are cited here in full, because they produced changes of measure and their materials are available: the sixty-third on 7 July, the seventieth on 15 October, and the seventy-first on 22 October.

On 7 July the joint body approved package insert revisions for both mRNA products, adding myocarditis and pericarditis to important precautions and directing that recipients be told in advance to seek medical attention if suggestive symptoms appeared. On 15 October it devoted a dedicated session to myocarditis-related events, recording that reported frequency was elevated in younger males. On 22 October, one week later, it determined that for males in their teens and twenties, given a clearly higher reported frequency after the Moderna product than the Pfizer product, the Pfizer product should also be selectable, subject to adequate information; the terms of that determination were circulated to paediatricians by the Japan Pediatric Society the following week.

Meetings between July and October are recorded as receiving myocarditis case listings, and did not produce changes of measure. The present author has not reconstructed the reasoning of each intervening meeting from its minutes, and does not assert that each expressly concluded that existing measures sufficed; what the published materials show is that the listings were tabled and that no measure changed. Even on that weaker reading the point holds. Reports accumulated and were repeatedly before the body; the measures in force did not change, and then they did. What is not recoverable from these materials is any rule which, published beforehand, would have indicated what response the accumulating figures were expected to produce, or at what point.

Two cautions constrain the use of this observation. The mature epidemiological picture was not available at the time: the English self-controlled case series reporting an incidence rate ratio of 11.76 for myocarditis after a second dose of mRNA-1273, and the Nordic cohort study reporting between nine and twenty-eight excess events per hundred thousand vaccinees after the same exposure in males aged sixteen to twenty-four, were published in 2022. Conduct must be assessed against what could be known at each date. And the absence of a rule is inferred from the materials examined, not from a systematic search.

Those later studies establish something relevant to design, however. Their stated conclusion has two parts: overall the risk of myocarditis was greater after SARS-CoV-2 infection than after vaccination, yet in younger men, particularly after a second dose of mRNA-1273, the ordering was reversed. A comparison drawn at population level does not transfer to the stratum in which vaccine-associated risk is highest and infection-associated risk lowest. Any rule operating on aggregate quantities would have been insensitive to exactly the finding that eventually moved the policy.

4. Why the mapping step may resist pre-specification

One explanation is worth stating, as a hypothesis compatible with the pattern above rather than a finding. Nothing in the design of this paper could establish it, and the proposal in Section 6 does not depend on it.

Under a discretionary arrangement, restricting an intervention previously authorised carries an implicit second message: that the earlier authorisation is being revisited. That invites questions directed at the decision-maker rather than the intervention — why the original

decision was made, why it was not revised earlier — and those questions are put to the institution that must make the revision. A reinforcing asymmetry concerns visibility: the costs of acting are concentrated, immediate and attributable, while the costs of not acting are diffuse, delayed and never definitively attributable, since the counterfactual is unobservable. No bad faith is required for such a comparison to tilt.

If the hypothesis holds, a published rule addresses part of it. Acting on a rule fixed beforehand is not an admission that the earlier authorisation was mistaken; it is the operation of something that formed part of that authorisation. The accountability question shifts from why the institution changed its mind to whether it did what it had published. It should be conceded that a rebuttable rule does not remove the difficulty entirely, and may in one respect increase it: a regulator departing from the presumption must now publish its reasons, which is an additional cost of explanation rather than a reduction. What the arrangement changes is which question the explanation must answer.

5. The distribution of risk and benefit

One further consideration applies specifically to preventive interventions, offered as a conditional design principle rather than a finding about any programme.

Benefit and risk are commonly described as weighed against each other, which suggests a single balance falling to one party. For an intervention administered across a population this need not hold: intervention-attributable risk falls on recipients, while benefit may accrue substantially to others. Both quantities are probabilistic, and the comparison between them is stratum-dependent, as Section 3.2 shows. What is distinctive is that two probability distributions are borne by partly different populations, which aggregate benefit-risk language conceals.

If a group is asked to accept risk for a benefit accruing substantially to others, something is owed to it: not exemption or compensation, but advance knowledge of the terms on which the risk is asked. That means publishing, before deployment, what response a given profile of observed harm would be expected to produce. Establishing that a particular group was in fact asked to bear risk primarily for others would require documenting the justification offered to it at each date and the direct benefit it could then expect; that work is not undertaken here, and the principle is stated conditionally.

6. What would have to be specified

6.1 Not a single magnitude

The most natural first thought — a threshold on the ratio of intervention-attributable to background incidence — is insufficient on its own, and stating why identifies what a rule must contain. A relative measure carries no information about absolute burden: a threefold elevation from one to three events per million is not the regulatory equivalent of a threefold elevation from one hundred to three hundred. No incidence measure distinguishes a transient and fully reversible event from a fatal one, or expresses how firmly the estimate is held, or captures the size of the exposed population, or the benefit forgone by restriction, which for a preventive product during an epidemic may be substantial and may itself vary by stratum.

These are, as Section 2.3 notes, the considerations European guidance already directs regulators to weigh. The proposal is not that they be identified but that their aggregation be stated: a published rule over the several dimensions, indicating what response a given profile is expected to produce.

6.2 Three candidate forms

Listing dimensions does not yield a rule. Some function must determine what profile produces what expectation, and its form is part of the proposal rather than a calibration detail. Three architectures are worth distinguishing.

The first is a set of severity-indexed thresholds: outcomes grouped into classes by seriousness and reversibility, each class carrying its own condition on excess risk, more serious classes carrying lower conditions. Within a class the rule reduces to a single comparison. This is the architecture of the drinking water tiers; it is transparent and matches how regulators already reason about seriousness. Its weaknesses are brittleness at class boundaries and dependence on a stable classification of outcomes, which for an emergent event may not exist.

The second is a conjunctive rule with lexicographic ordering: the presumption arises only where several conditions hold together, considered in a fixed order so that failure at an earlier stage ends the enquiry. This is the least susceptible to trading a weak showing on one dimension against a strong showing on another, but it is restrictive — a signal overwhelming on one dimension and marginally short on another produces no presumption at all.

The third is a weighted composite score with a stated cut-point: the most flexible and the most dangerous. Weights invite dispute, small changes can reverse conclusions, and a composite is the easiest of the three to influence through the measurement choices discussed in Section 6.5. If used, weights and cut-point must be published with the same status as the rule itself.

This paper takes no position among them beyond observing that the first is closest to arrangements already operating in an adjacent domain, and that the third carries risks requiring specific mitigation. What matters is that some such form must be chosen and published.

6.3 The rule is rebuttable, and departure must also be specified

The rule proposed here is not a trigger. Satisfaction of its conditions does not compel a regulatory measure. It raises a presumption of a defined response, from which a regulator may depart. A trigger would require that the rule be right, and the open questions in Section 6.5 make that a demanding requirement; a presumption lowers the cost of a mis-specified rule, because departure remains available. It does not remove the need to specify the rule well. A rule set too low would generate presumptions that are routinely departed from, which would discredit the arrangement; one set too high would raise none.

If departure is left unstructured, however, the arrangement fails in a predictable way. Discretion would move from the choice of standard to the choice of reasons for not applying it, and a regulator could satisfy the requirement indefinitely by writing that the criteria were met but that, on an overall view, no measure was warranted. Three requirements would prevent this without restoring a trigger.

The first is form. A departure must be published within the same period as the recommendation it accompanies, must state which conditions of the rule were met, and must be expressed as a departure from the rule rather than as a free-standing assessment. A statement that does not identify the presumption it declines to follow is not a departure but an evasion of the form.

The second is that grounds be enumerated. A rule should state in advance the categories of consideration capable of rebutting the presumption — for instance that the estimate is an artefact of a specified measurement choice, that a countervailing benefit not captured by the

rule applies to the affected stratum, or that a measure already in force achieves the same effect. The list need not be closed.

The third requirement is what makes the list workable and is the most important of the three. Where a regulator departs on a ground the rule did not enumerate, that ground must be added to the published rule prospectively. A novel reason may be used, but it may not be used silently, and it may not be used twice as though it were novel. This is what prevents the arrangement from decaying into the practice it is meant to replace: a standard that is quietly reconstructed each time it is applied. Under this requirement, the standard can still change, but only forward, and only in public.

6.4 Where such a rule would sit

The surrounding stages are already governed, and the rule proposed here does not duplicate them. Detection, notification, validation and assessment are timed and specified. The rule operates one stage later, on the content of the recommendation that assessment produces: given an assessed signal whose profile meets the published conditions, the presumption is a defined response, and the regulator may recommend otherwise on published grounds. The insertion point is the transition from assessment to recommendation, not the entry to assessment.

Two further constraints follow. Obligation must not be conditional on causal determination, or the most common ground for delay is restored. And where a required step does not occur within its window, the default should be procedural escalation rather than lapse — with the emphasis on procedural. A missed deadline caused by administrative delay or data interruption should advance the matter to the next stage of consideration and disclosure; it should not by itself produce a substantive restriction, since that would allow process failure rather than evidence to determine regulatory posture.

A rule applied to an estimate rather than a measurement also needs a decision rule under uncertainty, and that too is part of the architecture: whether comparison is against a point estimate or a confidence bound; whether one data source suffices or replication is required; how repeated looks at accumulating data are handled. On the last, the field is not starting from nothing, since sequential methods developed to control false positives under repeated testing are already in routine use in active surveillance.

6.5 Open questions

The central unresolved problem is what the rule should contain. Nuclear and drinking water limits were not derived from first principles either; they emerged from operating experience, negotiated standards and successive revision. That suggests a route but supplies no answer for a domain that has not begun.

Three questions follow. Who should set the rule, given that those with the relevant competence are often those whose conduct it would constrain? How should it be revised, given that a procedure too easily invoked defeats pre-commitment while one too rigid entrenches error — noting that the prospective-addition requirement in Section 6.3 supplies one route to revision but not a general one? And how should the rule handle a signal that is real while the relevant background rate is itself unstable, which is the hardest case technically and, during an epidemic, the likeliest?

A fourth question bears on whether the proposal can achieve what it claims. Any rule with institutional consequences creates an incentive to influence what is measured. Case definition, denominator, stratum boundaries, surveillance intensity and validation rules all affect whether conditions are met, and all involve judgment; for an emergent signal the relevant case definition may not exist as a specified category beforehand. Pre-commitment at the level of the rule may relocate discretion upstream into measurement. Two partial defences are available and would need development. The rule could specify a measurement hierarchy in advance, so that departures from the default specification are visible as departures. And where a regulator departs from the presumption on the ground that the estimate is a measurement artefact, the requirement in Section 6.3 would oblige it to publish the alternative specification and the result it yields — which converts a silent methodological choice into a disclosed one. Neither defence closes the problem, and it remains the strongest objection to the proposal.

These questions are why the conclusion here is limited to the proposition that a pre-specified escalation rule merits evaluation as an institutional option — not that it should be adopted, and not that adopting it would be anything other than difficult.

7. Objections

Four objections deserve notice, each answered in substance above. That existing arrangements already pre-specify a great deal is correct, and Section 2.3 sets out what they specify; the proposal adds one term, the rule connecting stated considerations to an expected outcome, and does so at the recommendation stage. That biological systems are too variable for numerical conditions has force against a claim that a rule could be derived precisely, but a rule raising a rebuttable presumption does not assert that values marginally below its conditions are safe; it states where the burden of explanation shifts. That the proposal would constrain regulatory flexibility is answered by its form: a regulator retains the power to reach any conclusion it considers correct, including that no measure is warranted, and loses only the ability to reach it without saying so against a published standard. And the mechanism proposed in Section 4 is not established, which is why the recommendation does not rest on it.

8. Conclusion

Post-market safety governance pre-specifies more than is often assumed. Outcomes, risk intervals and statistical thresholds are fixed before data are examined; notification and assessment run on deadlines measured in days; rationale is recorded and recommendations are published, including recommendations that nothing further is needed. Even the considerations bearing on those recommendations — severity, reversibility, strength of evidence, exposed population, vulnerable groups, the consequences of discontinuation — are set out in advance.

What was not identified in the instruments examined here is the rule that connects them. Which profile of considerations should issue in which response is settled while the signal is being assessed. Reasons are given and published, but a reader has no term stated beforehand against which to ask whether the conclusion departed from what had been undertaken.

The arrangement put forward for evaluation is correspondingly narrow: a published rule over several dimensions, raising a presumption of a defined response; departure permitted on published grounds expressed against the rule; and novel grounds of departure added to the rule prospectively, so that a standard may change but only forward and only in public. Nuclear operating limits and drinking water notification rules show that binding a stated condition to a stated consequence is institutionally achievable in complex domains. The 737 MAX and the Japanese myocarditis sequence show what the alternative looks like — in the

first, divergent national decisions with the certifying authority last and no common standard against which to assess the divergence; in the second, accumulating tabulations repeatedly before a committee, measures unchanged and then changed, with no rule stated beforehand indicating what response was expected or when.

Substantial questions remain: what the rule should contain, what form it should take, how statistical uncertainty enters it, and whether pre-commitment merely relocates discretion into measurement. Each is unresolved here, which is why the conclusion is evaluation rather than adoption.

The property that would make it worth evaluating is auditability, and it is worth being exact about which kind. Published reasons can already be read. What cannot presently be done is to ask whether a decision departed from what had been undertaken, because nothing was undertaken. A published rule supplies that term. It would not make regulators more likely to be right, and it would not remove their discretion. It would remove something narrower and, in the cases examined here, consequential: the freedom to settle the standard after the event it was to govern.

Declarations

Competing interests: The author declares no competing financial or non-financial interests.

Funding: This work received no external funding.

Ethics approval: Not required. The paper involves no human subjects, no animal experiments, and no primary data collection.

Data availability: No new datasets were generated. The paper draws on published literature and publicly available regulatory documents.

Peer review status: This is a working paper. It has not been peer reviewed.

Use of artificial intelligence: The author used Claude (Anthropic) for drafting assistance, structural revision, language refinement, and verification of bibliographic details against primary sources. All arguments and conclusions are the author’s own, and the author is responsible for the content.

References

All references were verified against primary or publisher sources during preparation of this paper.

Baggs, J., Gee, J., Lewis, E., Fowler, G., Benson, P., Lieu, T., et al. (2011). The Vaccine Safety Datalink: A model for monitoring immunization safety. Pediatrics, 127(Suppl 1), S45–S53. https://doi.org/10.1542/peds.2010-1722H

Elster, J. (2000). Ulysses unbound: Studies in rationality, precommitment, and constraints. Cambridge University Press.

European Medicines Agency. (2017). Guideline on good pharmacovigilance practices (GVP) — Module IX: Signal management (Rev 1). EMA/827661/2011 Rev 1. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-good-pharmacovigilance-practices-gvp-module-ix-signal-management-rev-1_en.pdf

European Union Aviation Safety Agency. (2019). Safety Directive 2019-01: Suspension of flight operations regarding Boeing 737-8 MAX and Boeing 737-9 MAX. Effective 12 March 2019, 19:00 UTC.

Federal Aviation Administration. (2020). Boeing 737 MAX timeline. https://www.faa.gov/sites/faa.gov/files/2022-08/737_MAX_Timeline.pdf

Japan Pediatric Society, Committee on Immunization and Infectious Diseases. (2021, October 29). Myocarditis-related events after COVID-19 vaccination: Information for paediatricians [in Japanese]. https://www.jpeds.or.jp/uploads/files/20211102_corona-1.pdf

Karlstad, Ø., Hovi, P., Husby, A., Hartläner, T., Pihlström, N., et al. (2022). SARS-CoV-2 vaccination and myocarditis in a Nordic cohort study of 23 million residents. JAMA Cardiology, 7(6), 600–612. https://doi.org/10.1001/jamacardio.2022.0583

Ministry of Health, Labour and Welfare of Japan. (2021, July 7). Revision of package inserts for Comirnaty intramuscular injection and COVID-19 Vaccine Moderna intramuscular injection. Material 2, 63rd meeting of the Subcommittee on Adverse Reactions, Health Sciences Council, and 12th meeting of the Safety Measures Investigation Committee, Pharmaceutical Affairs and Food Sanitation Council [in Japanese]. https://www.mhlw.go.jp/content/10601000/000802381.pdf

Ministry of Health, Labour and Welfare of Japan. (2021, August 4). Listing of myocarditis-related cases (reports from medical institutions). Material 1-5-2, 66th meeting of the Subcommittee on Adverse Reactions and 15th meeting of the Safety Measures Investigation Committee [in Japanese]. https://www.mhlw.go.jp/content/10601000/000816489.pdf

Ministry of Health, Labour and Welfare of Japan. (2021, October 15). Status of suspected myocarditis-related event reports. Material 1-1-1, 70th meeting of the Subcommittee on Adverse Reactions and 19th meeting of the Safety Measures Investigation Committee [in Japanese]. https://www.mhlw.go.jp/content/10601000/000844075.pdf

Patone, M., Handunnetthi, L., Saatci, D., Pan, J., Katikireddi, S. V., et al. (2022). Risk of myocarditis after sequential doses of COVID-19 vaccine and SARS-CoV-2 infection by age and sex. Circulation, 146(10), 743–754. https://doi.org/10.1161/CIRCULATIONAHA.122.059970

United States Code of Federal Regulations. Title 10, § 50.36 — Technical specifications. U.S. Nuclear Regulatory Commission.

United States Code of Federal Regulations. Title 40, Part 141, Subpart Q — Public notification of drinking water violations. U.S. Environmental Protection Agency. https://www.ecfr.gov/current/title-40/chapter-I/subchapter-D/part-141/subpart-Q

United States Nuclear Regulatory Commission. Operability guidance. https://www.nrc.gov/reactors/operating/licensing/techspecs/operability-guidance