



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Drug Safety Signals Gain Meaning Only When Exposure, Background Risk, and Time Are Reconciled

Ana Rodrigues^{1*}, Tiago Martins¹, Bruno Lopes²

¹Department of Pharmacovigilance and Drug Safety, Faculty of Pharmacy, University of Lisbon, Lisbon, Portugal.

²Department of Exposure–Risk Analysis, Faculty of Pharmacy, University of Porto, Porto, Portugal.

*Email: ana.rodrigues@ff.ulisboa.pt

ABSTRACT

Pharmacovigilance can detect unusual drug–event reporting patterns long before their clinical meaning is settled. The interpretive error is to treat signal strength, exposed-population risk, excess above expected occurrence, temporal plausibility, and causal credibility as though they were interchangeable quantities. They are not. Exposure denominators may be incomplete or misclassified; background event rates vary across populations, databases, phenotypes, and calendar periods; and the meaning of an observed excess depends on whether follow-up corresponds to a biologically credible risk period. Treatment switching, discontinuation, indication, and other time-dependent processes further complicate inference. Case-level dechallenge, rechallenge, and temporal coherence can strengthen or weaken an interpretation but cannot independently settle causality. This Current Opinion argues that drug safety signals should be interpreted through an explicitly proposed reconciliation of exposure credibility, background-event compatibility, biologically relevant time, dynamic treatment state, competing explanations, and case chronology. The framework is intended as a qualitative triage architecture, not a validated causality algorithm, prediction model, or regulatory standard.

Keywords: Pharmacovigilance, Disproportionality analysis, Exposure denominator, Background incidence, Risk window, Pharmacoepidemiology

Received: 15 December 2024; **Revised:** 20 January 2025; **Accepted:** 23 January 2025

INTRODUCTION

Pharmacovigilance increasingly operates across spontaneous reports, claims, electronic health records, registries, and other routinely collected data. These sources can complement one another, but they measure different aspects of drug safety and carry different ascertainment structures [1]. Signal communication itself reflects heterogeneous combinations of statistical, clinical, temporal, and contextual information rather than a single evidentiary standard [2]. Methods applied to routinely collected healthcare data are similarly diverse, with performance depending on the drug, event, data source, and design rather than on one universally preferred detection strategy [3].

The central interpretive problem therefore lies downstream of detection. A drug–event pair may be disproportionately reported without demonstrating that the event occurs more frequently among exposed patients than would otherwise be expected. Even when an apparent excess is observed, interpretation depends on who was actually exposed, how long exposure persisted, how the event was defined, what background incidence applies to the relevant population, and whether the event occurred within a biologically credible period. These distinctions become especially important when treatment changes over time or when the indication for treatment independently predicts the outcome.

This Current Opinion argues that signal detection and risk interpretation should be treated as analytically distinct stages. The article develops that position through exposure denominators, background-event rates, biologically relevant risk windows, time-varying treatment, confounding, and case-level temporal coherence. Their later integration into a proposed Time-Exposure-Background-Risk Triage Framework is deliberately qualitative: the aim is to discipline interpretation, not to manufacture a numerical causality score.

From disproportionality signal to risk interpretation

Disproportionality analysis asks whether a drug–event combination is reported more often than expected relative to a chosen reporting reference. That can be useful for screening, but signal prioritization need not depend on disproportionality alone. In prospective operational use, *vigiRank* combined several report characteristics and improved signal yield relative to the historical screening process in its evaluated setting [4]. The relevant lesson is not that a ranking algorithm estimates risk, but that statistical prominence and evidentiary priority are already separable concepts.

Interpretability also depends on analytical transparency. A meta-epidemiological assessment of published disproportionality studies identified substantial under-reporting of methodological choices, including the procedures used to select drug–event pairs and analytical thresholds [5]. A companion analysis found frequent interpretive spin and causal language in publications based on disproportionality methods [6]. These findings do not invalidate disproportionality analysis; they delimit what its outputs can support. A reporting odds ratio, for example, expresses a reporting contrast under a particular database and comparator structure, not an exposed-population incidence ratio.

Comparator choice is therefore part of the meaning of a signal rather than a technical afterthought. Different reference groups alter the reporting contrast and may introduce or mitigate biases related to therapeutic class, indication, co-medication, or reporting behavior [7]. A signal should consequently be read as conditional on its reference architecture.

The literature on published disproportionality signals also shows that investigators use case review, complementary data sources, clinical context, and other strategies to assess signal validity [8]. Those strategies should not be converted into an informal vote count in which more supporting sources automatically imply causality. Their value lies in testing different explanations.

Reporting standards provide another boundary. *READUS-PV* specifies information needed for transparent reporting of disproportionality analyses [9]. Compliance can improve reproducibility and appraisal, but it cannot transform a hypothesis-generating analysis into causal evidence.

Case-level information may nevertheless add something that automated screening misses. In the French pharmacovigilance database, incorporating causality-assessment features identified some safety information not detected through disproportionality approaches [10]. The appropriate inference is complementarity: individual chronology and clinical assessment may enrich a population-level signal, but neither positive rechallenge nor a high disproportionality statistic should be treated as self-sufficient proof.

Exposure denominators

A safety numerator without a defensible denominator cannot establish how commonly an event occurs among exposed patients. Yet exposure is rarely a simple binary variable. Dispensing may not equal ingestion; exposure occurring outside the observable system may be missed; and errors in capture or adherence can bias estimated associations in different directions depending on the underlying scenario [11].

Study entry also matters. New-user designs clarify treatment initiation, eligibility, and time zero, thereby avoiding some ambiguities produced when prevalent users and newly treated patients are mixed [12]. They do not eliminate residual confounding, but they make the population at risk more interpretable.

Linked claims and electronic health-record infrastructures can improve longitudinal coverage, provenance, and the range of available clinical information [13]. Greater data richness, however, should not be confused with automatic validity. The denominator remains conditional on whether treatment initiation, persistence, switching, adherence, and loss of observable follow-up are adequately captured.

Observed-to-expected analysis makes this dependence explicit because expected counts require exposure, a background incidence rate, and a defined risk period. An applied analysis of Bell's palsy after COVID-19 vaccination illustrates how different risk windows and stratification choices enter that calculation [14]. Such applications demonstrate the structure of the problem, not a universal solution for medicines with prolonged, intermittent, or poorly dated exposure.

Background event rates

A background rate is not a biological constant. In a large multinational study of adverse events relevant to vaccine safety, incidence varied substantially across databases and demographic strata [15]. The implication for drug safety is general: an expected count is meaningful only to the extent that the reference incidence describes a population, phenotype, and observation process compatible with the exposed group.

Historical comparison also introduces methodological risk. An empirical multidatabase evaluation showed that observed-to-expected approaches based on historical rates could retain substantial systematic error, with calibration and adjustment changing the balance between sensitivity and false-positive findings [16]. Therefore, a mathematically precise expected count may still be epidemiologically inappropriate.

Population structure matters within as well as across databases. Ontario background rates differed by event, age, and sex [17]. US estimates additionally varied across data sources and calendar periods, including pandemic-related changes for some outcomes [18]. These results caution against treating a pre-existing annual incidence estimate as indefinitely transportable.

Standardization helps but does not erase context. In the European ACCESS network, application of a common protocol across multiple databases still produced important between-database differences attributable to data capture and healthcare setting [19]. A larger multinational analysis likewise found cross-site variation in background incidence [20].

For pharmacovigilance, the practical question is therefore not simply “What is the background rate?” but “What background rate is defensible for this exposed population, this event definition, this healthcare setting, and this calendar period?” If that compatibility is weak, the uncertainty belongs inside the safety interpretation rather than being hidden behind a single expected value.

Biologically relevant risk windows

Time determines which person-time can plausibly contribute to a drug-related event. Observed-to-expected methods used during mass vaccination made this explicit because the interpretation changed with assumptions about exposure, expected incidence, and the duration of the post-exposure risk period [21]. Vaccination is an unusually tractable example because exposure dates are often well defined. Chronic treatment, intermittent administration, accumulation, withdrawal effects, delayed toxicity, and irreversible outcomes may require very different temporal reasoning.

Within-person designs demonstrate the same principle from another direction. In an application examining varenicline and cardiovascular outcomes, a self-controlled analysis produced a point estimate closer to an external benchmark than a nonuser comparison, although uncertainty remained substantial [22]. The finding supports the usefulness of within-person designs under appropriate assumptions; it does not establish their general superiority. The article therefore distinguishes a biologically relevant risk window from a merely convenient analytical window. The former should correspond as closely as available evidence allows to plausible latency, persistence of exposure effect, event onset, and outcome ascertainment. The latter may be defined primarily by data structure or convention. When those two windows diverge, an apparent excess may be diluted, displaced, or created by inclusion of person-time that is biologically irrelevant.

For this reason, risk-window appraisal is better treated as an adjudication problem than as a single calendar choice. The manuscript-derived states used below are explicitly proposed interpretive categories rather than validated classifications.

The analytical question that prose handles poorly is how differing signal sources, denominators, background rates, and temporal assumptions change what can legitimately be concluded from the same apparent safety concern. **Table 1** therefore adjudicates these dimensions side by side.

Table 1. Adjudicating whether observed safety events occupy a biologically defensible risk period

Proposed risk-window state	Signal source	Exposure denominator	Risk window	Event context and background risk	Principal competing bias mechanism	Decisive test or clarification	Permitted drug-safety endpoint
Well-dated acute exposure	Spontaneous report or linked longitudinal	Exposed persons or exposed person-time	Short interval consistent with	Background incidence matched to event	Reporting stimulation or transient	Re-estimate against compatible background	Temporally compatible excess requiring

	data with credible exposure date	can be defined with reasonable confidence	plausible onset	phenotype and comparable population	calendar effects	incidence and test alternative window definitions	further causal assessment
Persistent treatment with uncertain biological onset	Longitudinal healthcare data or accumulated case series	Person-time depends on persistence, refill capture, adherence, and discontinuation	Multiple plausible latency intervals	Baseline risk may evolve with disease severity and treatment duration	Exposure misclassification and depletion of susceptible patients	Evaluate alternative latency and persistence assumptions while updating exposure status	Association remains timing-sensitive rather than attributable
Poorly observed exposure timing	Spontaneous reports or incomplete dispensing history	Denominator incomplete or partly inferred	Nominal window exceeds confidence in actual exposure timing	Background rate may be estimable but cannot be aligned reliably to exposure	Misclassification of time at risk	Improve exposure ascertainment or restrict to cases with interpretable timing	Signal remains hypothesis-generating
Background-incompatible comparison	Any source with observed event count	Exposure may be adequate	Window may be plausible	Expected incidence comes from mismatched population, database, calendar period, or case definition	Transportability error in expected risk	Reconstruct expected incidence in a compatible population and observation system	Excess cannot be interpreted until expected risk is recalibrated
Recurrent or intermittent exposure	Prescription, administration, or case-level chronology	Denominator changes across exposed and unexposed intervals	Multiple recurrent windows	Background event hazard may itself vary over time	Time-varying confounding and carry-over	Model transitions between exposure states and test carry-over assumptions	Episode-specific association if temporal assumptions remain defensible
Delayed or persistent outcome	Longitudinal follow-up plus clinically characterized cases	Requires observation after exposure cessation	Window may extend beyond active treatment	Baseline incidence must remain comparable across extended follow-up	Informative censoring or arbitrary truncation of latency	Examine delayed onset, persistence, censoring, and alternative causal explanations	Delayed association remains plausible but unproven

The same adjudication can be represented visually as a lattice rather than a pipeline because exposure credibility, expected incidence, and temporal plausibility may agree or conflict independently. **Figure 1** makes the central interpretation boundary explicit.




Figure 1. The risk-time compatibility boundary separating report accumulation from interpretable excess risk

Time-varying exposure

Drug exposure can change after cohort entry through discontinuation, switching, dose modification, intermittent use, re-initiation, or loss of adherence. When analysis treats such exposure as fixed, the resulting person-time may no longer represent the treatment state that was clinically relevant when the event occurred.

Time-related biases in pharmacoepidemiology illustrate how strongly interpretation depends on alignment among eligibility, treatment, latency, and outcome ascertainment [23]. Different temporal errors can operate in different directions; they should therefore be diagnosed rather than invoked generically as a reason to distrust observational evidence.

Treatment switching creates an additional problem because post-baseline decisions may depend on prognosis, toxicity, effectiveness, or evolving disease. Analyses of switching and censoring show that estimates can differ according to how changing treatment status and post-treatment selection are handled [24]. No single longitudinal model is automatically correct. The analysis must correspond to the estimand, treatment trajectory, available covariates, and plausible causal structure.

This has a direct implication for safety signals. A denominator defined at treatment initiation may be appropriate for one question but misleading for another if substantial person-time accumulates after discontinuation or switching. Conversely, restricting follow-up too aggressively may miss delayed toxicity. Exposure should therefore be treated as a state that can transition over time whenever the pharmacology, clinical use pattern, or event latency makes those transitions consequential.

Table 2 translates these dynamic exposure states into proposed escalation decisions without implying that any state itself establishes causality or dictates regulatory action.

Table 2. Translating dynamic treatment states into interpretable safety-escalation questions

Time-varying exposure case	Signal source	Clinical seriousness	Preventability or reversibility	Cross-source concordance	Proposed escalation logic	Validation need	Permitted interpretive endpoint
Serious event concentrated during clearly exposed person-time	RWD with credible treatment dates, potentially supported by case reports	High clinical consequence	Depends on whether exposure can be stopped or modified	Concordance strengthens prioritization but is not required for interpretation	Prioritize focused causal evaluation if background risk and	Replicate with design addressing indication and time-varying confounding	Temporally concentrated association warranting intensified evaluation

timing are compatible							
Event occurs after discontinuation	Longitudinal data with reliable stop date	Depends on phenotype	Reversibility may be limited for delayed outcomes	Case chronology can clarify persistence after withdrawal	Escalate only if delayed effect remains biologically plausible	Test latency, carry-over, informative censoring, and alternative causes	Possible delayed association
Switching precedes the event	Claims, EHR, or treatment registry	Event-specific	Potentially preventable if switching reflects emerging toxicity	Source agreement may be difficult because exposure attribution changes	Reconstruct treatment sequence before attributing the event to one product	Model switching and post-treatment selection explicitly	Attribution remains treatment-state dependent
Repeated treatment episodes accompany repeated event timing	Administrative records plus case chronology	Event-specific	Potentially reversible between episodes	Recurrent timing may add coherence	Prioritize individual chronology and within-person evaluation where assumptions fit	Assess time-varying confounding and carry-over between episodes	Recurrent temporal association without automatic causal confirmation
Long-latency serious event	Longitudinal follow-up	Potentially high	Often incompletely reversible	Spontaneous reports may under-specify latency	Preserve extended follow-up rather than forcing an acute window	Validate exposure persistence, delayed biology, outcome definition, and comparator	Delayed signal requiring disease- and treatment-specific evaluation
Exposure history remains unstable or incompletely observed	SRS or fragmented RWD	Any seriousness	Usually indeterminate	Apparent concordance may reflect shared ascertainment error	Improve exposure reconstruction before increasing causal interpretation	Link data sources, test alternative exposure definitions, assess missingness	Signal remains insufficiently exposure-resolved

Confounding and indication

Safety associations are especially vulnerable when the clinical reason for prescribing a drug is itself associated with the event of interest. Confounding by indication can therefore create an apparent drug effect even when the treatment is not the cause, or conceal an effect when treated patients differ systematically from untreated patients. Active-comparator designs can improve clinical comparability by placing treated patients against patients with a similar therapeutic indication, but residual confounding can remain [25].

The active-comparator new-user design is consequently valuable but not universal. Suitable comparators may be unavailable, treatment patterns may differ materially, or the scientific question may concern a setting in which prevalent use, discontinuation, or alternative clinical strategies are more informative. Several comparator architectures can therefore be considered, each with context-dependent trade-offs [26].

For signal interpretation, this means that the background rate cannot be separated completely from the treatment-selection process. A population receiving a medicine may have a different underlying hazard of the event because of disease severity, comorbidity, prior treatment failure, surveillance intensity, or contraindication patterns. A numerical excess over an inadequately matched reference population can therefore remain confounded even when exposure and event timing are measured accurately.

The proposed triage logic developed later does not treat confounding as a final adjustment step. It treats competing explanations as part of the meaning of the signal itself. Exposure credibility, compatible expected incidence, and plausible timing can make an association more interpretable, but they do not remove indication-related differences or establish causality. Comparator choice and clinical context must remain visible when a signal is escalated from detection toward risk interpretation.

Dechallenge, rechallenge, and temporal coherence

Case chronology can change the interpretation of a population-level signal because it addresses a different question: whether exposure, event onset, withdrawal, recovery, and re-exposure form a sequence compatible with a drug-related effect. Reviews of communicated pharmacovigilance signals show that temporal information, including dechallenge and rechallenge, is frequently considered during signal assessment [2]. In the French pharmacovigilance database, individual causality-assessment information also identified some safety information that automated disproportionality approaches did not recover [10]. These observations support complementarity, not a hierarchy in which case chronology overrides population evidence.

Dechallenge is most informative when an event is expected to improve after exposure ceases and competing treatments, disease progression, and supportive interventions do not offer equally plausible explanations. Rechallenge can be more persuasive when a comparable event recurs after re-exposure, but it may be ethically unavailable, clinically undesirable, incompletely documented, or difficult to interpret when background recurrence is common. Published strategies for increasing the validity of disproportionality signals therefore use case assessment alongside complementary data and contextual evidence rather than treating one feature as decisive [8]. Positive rechallenge should consequently be treated as an evidentiary contribution whose weight depends on phenotype, timing, alternative causes, and exposure certainty [10].

Negative or absent dechallenge is less symmetric. Failure to recover after withdrawal may reflect irreversible injury, prolonged biological persistence, delayed pathophysiology, or an outcome whose course is poorly coupled to drug elimination. Conversely, improvement after withdrawal may coincide with treatment of the event or spontaneous resolution. Missing rechallenge usually means that an informative observation does not exist, not that evidence against causality has been generated [2]. Temporal coherence should therefore be interpreted as a pattern rather than a binary causality test [8].

The proposed temporal-state interpretation below preserves what is learned at each stage of the exposure–event sequence while preventing case chronology from becoming detached from denominator quality, risk-window plausibility, and competing explanations.

Table 3. Temporal states governing how dechallenge and rechallenge can modify safety-signal meaning

Proposed temporal state	Entry condition	Signal source and exposure denominator	Risk-window relation	State-specific evidence and event context	Transition or persistence	Time-dependent validation need	Proposed framework endpoint
Interpretable onset during exposure	Exposure timing and event onset are credible	Case chronology with treatment records where available	Onset occurs within a defensible latency interval	Phenotype is characterized and immediate alternatives are assessed	Event persists, resolves, or evolves after exposure change	Confirm exposure timing, competing causes, and expected natural history	Temporally coherent evidence that can strengthen further evaluation
Improvement after withdrawal	Treatment stops before outcome resolution	Individual case plus observable treatment history	Recovery timing is compatible with plausible effect decay	Improvement is not clearly explained by another intervention or spontaneous course	Resolution follows withdrawal	Assess reversibility, concomitant treatment, and recovery kinetics	Supportive dechallenge evidence without causal confirmation
Persistence after withdrawal	Event continues	Longitudinal case or	Persistence remains compatible	Outcome biology permits	Symptoms or	Distinguish persistent drug effect from	Dechallenge remains indeterminate

	after cessation	healthcare data	with delayed or irreversible effect	continued injury or slow recovery	findings continue	disease progression and unrelated persistence	rather than automatically negative
Recurrence after re- exposure	Re-exposure follows improvement or a prior event	Case chronology with credible renewed exposure	Recurrence occurs within a comparable plausible interval	Phenotype resembles the earlier event and background recurrence is considered	Event reappears after renewed exposure	Verify exposure, event similarity, alternative triggers, and recurrence propensity	Stronger temporal coherence requiring population contextualization
No re- exposure	Treatment is not restarted	Case report or clinical record	Rechallenge window does not exist	Clinical decision may reflect seriousness, uncertainty, or alternatives	No subsequent exposure state	Do not infer evidence for or against causality from missing rechallenge	Unresolved temporal evidence
Discordant chronology	Event precedes exposure or occurs outside a defensible temporal model	Source with sufficiently reliable dates	Timing conflicts with the suspected drug–event relation	Alternative explanations become comparatively stronger	Sequence remains inconsistent	Reassess event definition, exposure dates, competing causes, and temporal assumptions	Interpretation weakened unless another defensible temporal model exists

The relationship between case chronology and population-level evidence is nonlinear: a convincing sequence can strengthen a plausible signal, whereas incompatible timing can weaken it, yet neither should erase denominator or background-risk deficiencies. **Figure 2** therefore places temporal coherence inside a causal-bias diagnostic map rather than at the end of a one-way causality pathway.




Figure 2. Temporal coherence gates that prevent case chronology from bypassing population-level bias

Caption. Case chronology, population-level occurrence, and causal-design information remain analytically separate. Dechallenge and rechallenge enter a temporal-coherence field that is cross-checked against exposure certainty, biologically credible timing, background-event compatibility, treatment switching, and competing

clinical explanations. Solid connectors denote evidence-supported dependencies; dashed connectors denote the article's proposed conditional integration; dotted connectors denote unresolved or definition-sensitive relationships. Validation gates prevent case-level coherence from becoming causal proof and prevent population-level excess from bypassing temporal and bias assessment.

Alt text. Wide left-to-right causal-bias diagnostic map with three horizontal lanes. The upper lane contains case chronology, including onset, withdrawal, recovery, persistence, and re-exposure. The middle lane contains exposure denominator, background-event compatibility, and risk-window information. The lower lane contains confounding by indication, comparator choice, treatment switching, and time-related bias. Cross-links converge on central validation gates for temporal coherence, competing explanations, and cross-source compatibility. Beyond the gates are proposed strengthened, weakened, and unresolved signal states, explicitly separated from causal confirmation.

A time-exposure–background-risk triage framework

The evidence reviewed here supports a distinction between detecting unusual reporting and interpreting clinically meaningful excess risk. Published signal-validity strategies, observed-to-expected applications, and methodological experience with risk periods show that contextual information changes what a signal can support [8, 14, 21]. The proposed Time-Exposure–Background-Risk Triage Framework organizes that contextualization without numerical weights.

Its first question is whether exposure is resolved sufficiently to define meaningful person-time. The second asks whether expected event occurrence is compatible with the exposed population, phenotype, data source, and calendar period. The third asks whether events occupy a biologically defensible risk period rather than an arbitrary analytical interval. These are not sequential hurdles. Weakness in one dimension changes how evidence in the others can be interpreted.

The next layer asks whether treatment and causal structure remain stable during follow-up. Time-related bias, switching or censoring, and confounding by indication can each alter an apparent drug–event association even when exposure dates and outcomes are captured [23–25]. Dynamic treatment state and competing explanations are therefore interpretation channels rather than adjustments added after a signal has already been judged credible. Case chronology then modifies, rather than determines, interpretation. Coherent dechallenge or rechallenge may strengthen concern when denominator, background-risk, and timing assumptions are already defensible. Discordant or uninterpretable chronology may weaken or leave the signal unresolved. Proposed outputs are consequently qualitative states such as sufficiently contextualized for focused causal evaluation, materially weakened by competing evidence, or unresolved because essential inputs remain incompatible. They are not probabilities of causation, regulatory grades, or predictions of future harm.

Regulatory implications

Regulatory usefulness depends on more than statistical prominence. Routinely collected healthcare data can add evidence, but analytical method, reporting transparency, provenance, and completeness must remain visible in interpretation [3, 9, 13]. A signal that cannot specify exposure, event definition, comparator structure, and analytic time should not acquire certainty merely because similar associations appear in several data sources.

Observed-to-expected methods can provide rapid contextualization, but historical comparisons can retain systematic error [16]. Experience from mass vaccination similarly shows that expected rates and risk periods are assumption dependent [21]. Regulatory interpretation should therefore distinguish numerical elevation from an excess whose denominator, background reference, and temporal architecture are sufficiently compatible to justify focused causal investigation.

Confounding requires the same discipline. Active comparators can reduce confounding by indication in suitable settings [25], while alternative comparator designs may be preferable when clinical use patterns make an active-comparator new-user design unsuitable [26]. The proposed framework does not prescribe one design. Its potential value lies in making the reasons for escalation, downgrading, or continued uncertainty explicit.

An unresolved state is therefore legitimate. Some signals may remain important enough to monitor despite incompatible background data, uncertain exposure, or incomplete chronology; others may justify targeted observational studies, improved case collection, or cross-source replication. Future evaluation of the proposed framework should test whether it improves transparency and consistency of reasoning rather than assuming that more affirmative safety conclusions indicate better performance.

CONCLUSION

Drug safety signals become interpretable only when the quantities beneath them remain distinct. Reporting disproportionality identifies unusual reporting, exposure denominators define who and when was actually at risk, background rates define what would otherwise be expected, and biologically credible timing determines whether observed person-time is relevant to the suspected effect. Dynamic treatment, confounding, and case chronology can then strengthen, weaken, or leave that interpretation unresolved.

The proposed Time-Exposure-Background-Risk Triage Framework is a qualitative discipline for reconciling these dimensions, not a validated causality algorithm. Its premise is deliberately bounded: stronger safety interpretation requires better alignment among the signal observed, the population actually exposed, the risk expected without the drug, and the time in which a drug-related event could plausibly occur.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Trifirò G, Sultana J, Bate A. From big data to smart data for pharmacovigilance: The role of healthcare databases and other emerging sources. *Drug Saf.* 2018;41(2):143-9. doi:10.1007/s40264-017-0592-4
2. Sartori D, Aronson JK, Norén GN, Onakpoya IJ. Signals of adverse drug reactions communicated by pharmacovigilance stakeholders: A scoping review of the global literature. *Drug Saf.* 2023;46(2):109-20. doi:10.1007/s40264-022-01258-0
3. Coste A, Wong A, Bokern M, Bate A, Douglas IJ. Methods for drug safety signal detection using routinely collected observational electronic health care data: A systematic review. *Pharmacoepidemiol Drug Saf.* 2023;32(1):28-43. doi:10.1002/pds.5548
4. Caster O, Sandberg L, Bergvall T, Watson S, Norén GN. vigiRank for statistical signal detection in pharmacovigilance: First results from prospective real-world use. *Pharmacoepidemiol Drug Saf.* 2017;26(8):1006-10. doi:10.1002/pds.4247
5. Khouri C, Revol B, Lepelley M, Mouffak A, Bernardeau C, Salvo F, et al. A meta-epidemiological study found lack of transparency and poor reporting of disproportionality analyses for signal detection in pharmacovigilance databases. *J Clin Epidemiol.* 2021;139:191-8. doi:10.1016/j.jclinepi.2021.07.014
6. Mouffak A, Lepelley M, Revol B, Bernardeau C, Salvo F, Pariente A, et al. High prevalence of spin was found in pharmacovigilance studies using disproportionality analyses to detect safety signals: A meta-epidemiological study. *J Clin Epidemiol.* 2021;138:73-9. doi:10.1016/j.jclinepi.2021.06.022
7. Gravel CA, Douros A. Considerations on the use of different comparators in pharmacovigilance: A methodological review. *Br J Clin Pharmacol.* 2023;89(9):2671-6. doi:10.1111/bcp.15802
8. Fusaroli M, Salvo F, Bernardeau C, Idris M, Dolladille C, Pariente A, et al. Mapping strategies to assess and increase the validity of published disproportionality signals: A meta-research study. *Drug Saf.* 2023;46(9):857-66. doi:10.1007/s40264-023-01329-w
9. Fusaroli M, Salvo F, Begaud B, AlShammari TM, Bate A, Battini V, et al. The reporting of a disproportionality analysis for drug safety signal detection using individual case safety reports in pharmacovigilance (READUS-PV): Development and statement. *Drug Saf.* 2024;47(6):575-84. doi:10.1007/s40264-024-01421-9
10. Berbain T, Pariente A, Miremont-Salamé G, Grandvuillemin A, Micallef J, Chouchana L, et al. Contribution of causality assessment for an automated detection of safety signals: An example using the French pharmacovigilance database. *Drug Saf.* 2020;43(3):243-53. doi:10.1007/s40264-019-00887-2
11. Hempenius M, Groenwold RHH, de Boer A, Klungel OH, Gardarsdottir H. Drug exposure misclassification in pharmacoepidemiology: Sources and relative impact. *Pharmacoepidemiol Drug Saf.* 2021;30(12):1703-15. doi:10.1002/pds.5346

12. Her QL, Rouette J, Young JC, Webster-Clark M, Tazare J. Core concepts in pharmacoepidemiology: New-user designs. *Pharmacoepidemiol Drug Saf.* 2024;33(12). doi:10.1002/pds.70048
13. Desai RJ, Marsolo K, Smith J, Carrell D, Penfold R, Pillai HS, et al. The FDA Sentinel Real World Evidence Data Enterprise (RWE-DE). *Pharmacoepidemiol Drug Saf.* 2024;33(10). doi:10.1002/pds.70028
14. van der Boom MDX, van Eekeren R, van Hunsel FPAM. Observed-over-expected analysis as additional method for pharmacovigilance signal detection in large-scaled spontaneous adverse event reporting. *Pharmacoepidemiol Drug Saf.* 2023;32(7):783-94. doi:10.1002/pds.5610
15. Li X, Ostropolets A, Makadia R, Shoaibi A, Rao G, Sena AG, et al. Characterising the background incidence rates of adverse events of special interest for covid-19 vaccines in eight countries: Multinational network cohort study. *BMJ.* 2021;373. doi:10.1136/bmj.n1435
16. Li X, Lai LYH, Ostropolets A, Arshad F, Tan EH, Casajust P, et al. Bias, precision and timeliness of historical (background) rate comparison methods for vaccine safety monitoring: An empirical multi-database analysis. *Front Pharmacol.* 2021;12:773875. doi:10.3389/fphar.2021.773875
17. Nasreen S, Calzavara A, Buchan SA, Thampi N, Johnson C, Wilson SE, et al. Background incidence rates of adverse events of special interest related to COVID-19 vaccines in Ontario, Canada, 2015 to 2020, to inform COVID-19 vaccine safety surveillance. *Vaccine.* 2022;40(24):3305-12. doi:10.1016/j.vaccine.2022.04.065
18. Moll K, Lufkin B, Fingar KR, Zhou CK, Tworowski E, Shi C, et al. Background rates of adverse events of special interest for COVID-19 vaccine safety monitoring in the United States, 2019-2020. *Vaccine.* 2023;41(2):333-53. doi:10.1016/j.vaccine.2022.11.003
19. Willame C, Dodd C, Durán CE, Elbers RJHJ, Gini R, Bartolini C, et al. Background rates of 41 adverse events of special interest for COVID-19 vaccines in 10 European healthcare databases - an ACCESS cohort study. *Vaccine.* 2023;41(1):251-62. doi:10.1016/j.vaccine.2022.11.031
20. Phillips A, Jiang Y, Walsh D, Andrews N, Artama M, Clothier H, et al. Background rates of adverse events of special interest for COVID-19 vaccines: A multinational Global Vaccine Data Network (GVDN) analysis. *Vaccine.* 2023;41(42):6227-38. doi:10.1016/j.vaccine.2023.08.079
21. Gordillo-Marañón M, Candore G, Hedenmalm K, Browne K, Flynn R, Piccolo L, et al. Lessons learned on observed-to-expected analysis using spontaneous reports during mass vaccination. *Drug Saf.* 2024;47(7):607-15. doi:10.1007/s40264-024-01422-8
22. Yokoyama R, Takeuchi Y, Kumamaru H, Matsuyama Y. The utility of the self-controlled study design for pharmacoepidemiological studies without an active comparator medication using a medical information database: An application to assess the risk of varenicline on cardiovascular outcomes. *Pharmacoepidemiol Drug Saf.* 2023;32(10):1068-76. doi:10.1002/pds.5634
23. Suissa S, Dell'Aniello S. Time-related biases in pharmacoepidemiology. *Pharmacoepidemiol Drug Saf.* 2020;29(9):1101-10. doi:10.1002/pds.5083
24. Belviso N, Zhang Y, Aronow HD, Wyss R, Barbour M, Kogut S, et al. Addressing posttreatment selection bias in comparative effectiveness research, using real-world data and simulation. *Am J Epidemiol.* 2022;191(2):331-40. doi:10.1093/aje/kwab242
25. Sendor R, Stürmer T. Core concepts in pharmacoepidemiology: Confounding by indication and the role of active comparators. *Pharmacoepidemiol Drug Saf.* 2022;31(3):261-9. doi:10.1002/pds.5407
26. Wintzell V, Svanström H, Pasternak B. Selection of comparator group in observational drug safety studies: Alternatives to the active comparator new user design. *Epidemiology.* 2022;33(5):707-14. doi:10.1097/EDE.0000000000001521