Annals of Pharmacy Practice and Pharmacotherapy
ISSN: 3062-4436
2026, Volume 6, Issue 1, Page No: 41-50

Available online at: www.galaxypub.co/page/journals

Galaxy Publication

From Alert Burden to Preventable Harm: A Risk-Adaptive Medication-Safety Model for Interruptive Alerts, Silent Surveillance, Pharmacist Review, and Clinical Escalation

João Silva[1]*, Pedro Costa[1], Ana Beatriz[2]

[1] Department of Medication Safety and Clinical Alerts, Faculty of Pharmacy, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil.
[2] Department of Risk-Adaptive Safety Models, Faculty of Pharmacy, University of Campinas, Campinas, Brazil.

***E-mail** ✉ joao.silva@ufrj.br

Received: 19 November 2025; **Revised:** 09 February 2026; **Accepted:** 12 February 2026

ABSTRACT

Medication-related clinical decision support is commonly judged through alert counts, override rates, acceptance rates, or rule-level technical performance. These measures describe system activity but do not directly establish whether preventable medication harm has been reduced. This Current Opinion argues that medication-safety signals should instead be assigned to response modes according to the combination of potential harm, immediacy, contextual uncertainty, persistence, interpretive requirements, and responsibility for clinical action. Evidence from interruptive alerts, noninterruptive decision support, pharmacist-mediated review, and post-implementation optimization suggests that no single delivery mode is appropriate for all medication risks. Highly consequential and time-sensitive hazards may justify interruption when signals are sufficiently specific, whereas lower-immediacy or context-dependent signals may be better managed through silent surveillance followed by contextual reassessment. Pharmacist review can add an interpretive layer when rule correctness does not establish patient relevance, and escalation requires explicit transfer of responsibility rather than simple forwarding of another notification. We therefore propose a risk-adaptive medication-safety model in which signals can migrate among surveillance, interruption, pharmacist review, and clinical escalation as risk and uncertainty change. The model is conceptual rather than validated. Its usefulness will depend on local data quality, staffing, workflow, patient heterogeneity, system configuration, and prospective evaluation using patient-safety and balancing outcomes rather than alert-volume metrics alone.

Keywords: Clinical decision support, Medication safety, Alert fatigue, Preventable harm, Pharmacist review, Clinical escalation

How to Cite This Article: Silva J, Costa P, Beatriz A. From Alert Burden to Preventable Harm: A Risk-Adaptive Medication-Safety Model for Interruptive Alerts, Silent Surveillance, Pharmacist Review, and Clinical Escalation. Ann Pharm Pract Pharmacother. 2026;6(1):41-50. https://doi.org/10.51847/5CpTfj2hQs

Introduction

Medication-safety decision support has often been optimized around a visible operational problem: too many alerts. Yet “alert burden” is not a single clinically meaningful construct. A recent systematic review found substantial heterogeneity in how alert fatigue is defined and measured, with alert quantity, override, acceptance, and related proxies frequently used without demonstrating that they capture sustained deterioration in appropriate clinical response [1]. Consequently, a medication-safety system can appear improved because fewer notifications are displayed while remaining unchanged—or even worsened—in its ability to prevent consequential harm.

Technological optimization does not remove this measurement problem. Artificial-intelligence approaches are increasingly being applied to medication-alert prioritization and refinement, but implementation experience and external validation remain limited relative to technical development [2]. The central question is therefore not whether an algorithm can rank, suppress, or personalize alerts, but whether the resulting delivery strategy preserves clinically important detection while allocating attention in a way that clinicians and health systems can safely act upon.

Patient vulnerability further complicates that judgment. In older people, medication-related clinical decision support has shown more consistent evidence for changes in prescribing processes than for reductions in adverse clinical outcomes, underscoring the difference between improving a proximal decision and demonstrating prevention of patient harm [3]. The same distinction applies to interaction detection. Potential drug–drug interactions represent possible hazards, whereas clinically manifested harm depends on exposure, susceptibility, monitoring, treatment context, and response; multicentre evidence illustrates why the prevalence of detected interactions should not be equated with actual harmful events [4].

This Current Opinion therefore reframes medication-alert optimization around preventable-harm risk rather than alert volume. The argument distinguishes four response modes: interruptive alerts for selected time-critical or high-severity risks; silent surveillance for signals whose meaning depends on persistence or accumulating context; pharmacist review where patient-level interpretation is needed; and clinical escalation when responsibility must transfer to an actor capable of timely intervention. The integrated architecture advanced here is proposed rather than validated. Its purpose is to make the routing decision explicit: not merely whether a signal should exist, but who should see it, when they should see it, under what degree of interruption, with what information, and with what obligation to act.

Why alert volume is a poor proxy for medication-safety value

A low-volume alerting system is not necessarily a high-value one. National evaluation of hospital medication-related decision support has demonstrated a practical trade-off between reducing unnecessary alerting and preserving detection of unsafe medication orders [5]. The implication is straightforward: indiscriminate suppression can improve burden metrics while sacrificing clinically relevant sensitivity. Alert optimization therefore requires a value function broader than "fewer alerts."

The performance of an alert also depends on the environment into which it is deployed. After migration to a commercial electronic health record, interruptive drug–drug interaction alerts in one health system became less effective, illustrating that identical clinical concepts can behave differently after changes in configuration, interface, or workflow [6]. Alert modality should therefore not be treated as an intrinsic property of a drug interaction rule; its practical value emerges from the interaction among rule logic, display design, clinical context, and user response.

Patient outcomes provide an even stricter test. A recent systematic review found that electronic drug–drug interaction alerts can influence clinician behavior, but evidence for meaningful patient benefit and prevention of targeted interaction-related harm remains uncertain [7]. Process improvement is therefore necessary but insufficient evidence of medication-safety value.

Conversely, burden reduction can be defensible when it is selective and evaluated against loss of detection. Customized drug–drug interaction screening intervals reduced alert incidence and false-positive burden in one hospital without an observed reduction in sensitivity [8]. Such findings support targeted suppression, not suppression as an end in itself. The relevant question is whether low-value repetition can be removed while the system remains capable of identifying clinically consequential states.

Interruptive alerts for time-critical and high-severity risk

Interruptive alerts impose an immediate cognitive and workflow cost, so their justification should be correspondingly demanding. Hard-stop alerts can change clinician behavior, but systematic review evidence also documents unintended consequences, including workflow disruption and opportunities for workarounds [9]. Interruption is therefore best understood as a scarce safety intervention rather than the default display mode for every technically valid signal.

Evidence from intensive care provides a stronger case for selective interruption. In a cluster-randomized stepped-wedge trial across nine intensive care units, clinically tailored alerts targeting a restricted set of high-risk drug combinations reduced administration of those combinations [10]. The trial does not establish that every severe interaction should trigger an interruptive alert, nor does its process outcome prove universal reduction of adverse drug events. It does, however, support the principle that a narrowly selected, clinically contextualized set of high-risk signals can justify prominent interruption.

Severity alone is still insufficient. Evaluation of anticoagulant-related decision support in Dutch hospital pharmacies showed that even within a medication domain conventionally regarded as high risk, generated alerts

vary in relevance and operational burden [11]. A label such as "anticoagulant" or "high-alert medication" cannot substitute for evaluating the particular hazard, timing, patient state, and information needed for action.

This boundary is reinforced by evidence that context-insensitive drug–drug interaction alerting may fail to reduce interaction-related harm in hospitalized patients [12]. The proposed criterion is therefore conjunctive: interruption should require not simply theoretical severity, but a credible combination of **severity**, immediacy, preventability, sufficient contextual specificity, and an action that can reasonably be taken at the moment of interruption.

Table 1 operationalizes this proposed justification logic without assigning numerical thresholds.

Table 1. When a medication hazard warrants interruptive alerting

Alertable hazard	Immediacy	Severity	Preventability	Information needed at interruption	False-positive cost	Escalation destination
High-risk drug combination with an actionable ordering decision	Immediate or near-immediate	Potentially serious	Action possible before administration	Current drugs, dose, timing, relevant patient context	Workflow disruption; habituation if poorly targeted	Ordering clinician or responsible clinical team
Potential interaction whose harm depends on patient-specific conditions	Context dependent	Potentially serious	Conditional	Relevant laboratory results, organ function, exposure context, monitoring status	Unnecessary interruption if context is absent	Pharmacist or clinician when context remains uncertain
Hard-stop candidate in which proceeding could create an immediate safety hazard	Immediate	High	Action must occur before continuation	Minimum information sufficient to establish that blocking is justified	Delay, workaround, or abandonment of appropriate therapy	Responsible prescriber with defined override/escalation route
High-alert medication signal with uncertain patient relevance	Variable	Potentially high	Often preventable if correctly interpreted	Indication, co-medication, dose, duplication, patient susceptibility	High burden if medication class alone drives alerting	Pharmacist review before broader clinical escalation

Note: Entries represent the article's proposed interruptive-alert justification register derived from the evidence and argument developed above. They are qualitative decision states, not validated thresholds or universally applicable routing rules.

Silent surveillance for lower-immediacy or context-dependent signals

Some medication-safety signals should remain visible to the safety system without forcing an immediate interruption. Noninterruptive medication decision support is clinically feasible: a noninterruptive alert increased take-home naloxone prescribing in an emergency-department setting, while observational work with a passive medication-related decision-support system shows that alerts can be presented without requiring clinicians to stop and respond synchronously [13, 14]. These studies establish the viability of noninterruptive delivery in specific contexts; they do not demonstrate that silent approaches are universally safer or more effective than interruptive alerts.

Silent surveillance becomes more useful when a signal has temporal meaning. Longitudinal event analysis of passive medication alerts has shown that alert persistence and disappearance can be tracked across evolving medication states [15]. However, disappearance of an alert cannot be assumed to mean that the alert caused a corrective action. A prescription may change for unrelated clinical reasons, and a persistent alert may remain appropriate if the relevant risk is accepted and monitored.

For that reason, silent surveillance should not be equated with suppression. It is better conceptualized as an active monitoring state in which the system retains a signal, refreshes its context, and determines whether persistence, worsening severity, new patient information, or failure of expected follow-up warrants another response mode. Nor should movement be unidirectional. Optimization of interruptive decision support for antithrombotic duplicate prescribing demonstrates that a narrowly defined signal can remain appropriately interruptive after refinement [16]. Some signals should therefore migrate upward toward interruption, while others may move downward after contextual resolution.

Table 2 translates these principles into a proposed surveillance-to-review routing structure.

Table 2. Routing medication-safety signals from silent surveillance to active review

Surveillance signal	Persistence requirement	Contextual data needed	Pharmacist review trigger	Migration-to-interruptive criteria	Closure logic
Preventive medication opportunity without immediate danger	Reassess while eligibility persists	Indication, prior treatment, current encounter context	Clinical ambiguity or repeated non-action	Risk becomes time-sensitive or immediate action becomes necessary	Action completed, eligibility resolves, or responsibility documented
Passive medication alert of uncertain patient relevance	Persist until relevant context becomes available or risk resolves	Medication history, dose, laboratory and clinical state	Automated rule is technically valid but patient relevance remains unclear	Context demonstrates a high-severity, actionable state	Pharmacist/clinician adjudicates relevance or surveillance condition clears
Recurrent or persistent alert across medication changes	Persistence itself becomes informative	Longitudinal medication changes and related clinical data	Signal remains unresolved despite repeated reassessment	Persistence combines with increasing severity or narrowing uncertainty	Risk resolves, accepted monitoring plan is documented, or signal is escalated
Signal previously handled noninterruptively but now meeting higher-risk conditions	No fixed duration; migration is state dependent	Updated exposure, patient susceptibility and immediacy	Human interpretation needed before increasing alert intensity	Current state justifies immediate workflow disruption	Migrate to interruption, clinical escalation, or return to surveillance after resolution

Note: The routing states are proposed. "Persistence" does not imply causality, and migration should depend on updated clinical context rather than on elapsed time alone.

Pharmacist review as an interpretive safety layer

Automated medication alerts answer a narrow question: whether the programmed conditions for a rule have been met. They do not necessarily answer the clinically harder question of whether the alert is relevant to this patient at this moment. Clinical-pharmaceutical assessment of medication decision-support alerts has demonstrated a meaningful distinction between content appropriateness and patient relevance, supporting the need for a human interpretive step for selected signals [17].

That role is particularly apparent when appropriateness depends on indication, duration, transition status, or information that is difficult to encode in a simple rule. At hospital discharge, pharmacist-evaluated decision support has been used to assess potentially missing or inappropriate proton-pump inhibitor prescribing, illustrating how an automated signal can function as the starting point for contextual medication review rather than as a complete recommendation [18].

Pharmacist mediation can also reduce the number of technically valid but clinically low-relevance signals forwarded to prescribers. In one implementation, clinical pharmacists reviewed medication-CDSS alerts for technical validity and pharmaceutical relevance before selected transmission to physicians [19]. This supports a layered architecture in which automated detection and clinician interruption are not the only possible endpoints.

Prospective implementation of real-time pharmaceutical algorithms targeting high-alert medications similarly separated algorithmic detection, pharmacist characterization of drug-related problems, and subsequent pharmacist intervention [20]. The positive predictive value and usefulness of such systems therefore depend on more than the rule engine itself.

The model proposed here does not assume that every signal requires pharmacist review, nor that pharmacists are the only professionals capable of contextual adjudication. Rather, pharmacist review is positioned as a selective interpretive safety layer for signals whose potential consequence is meaningful but whose actionability remains uncertain after automated surveillance. Its feasibility depends on staffing, expertise, information access, prioritization, and a clear route onward when interpretation indicates that clinical action is required.

Clinical escalation and responsibility transfer

A medication-safety signal has little practical value if it reaches no person who can resolve it. Escalation should therefore mean more than increasing visual salience or forwarding an alert. It should transfer a defined problem to a named actor with sufficient information, authority, and expectation of follow-through. In the URGEIM randomized clinical trial, a structured pharmacist intervention after medication-related emergency-department

events improved selected downstream outcomes, supporting the importance of active follow-up once a medication-related problem has been identified [21]. The finding does not validate a universal pharmacist-led escalation pathway, but it shows why detection and responsibility transfer should be treated as separate stages.

Transition-of-care evidence reinforces this distinction. Pharmacist-led medication reconciliation at hospital discharge reduced medication errors, while short-term healthcare utilization did not clearly improve [22]. A closed loop can therefore succeed at the medication-process level without guaranteeing broader clinical or health-service effects. Escalation metrics should accordingly distinguish acknowledgement, clinical action, resolution, and downstream patient outcomes.

Responsibility also requires organizational stewardship. The Clickbusters initiative demonstrated how multidisciplinary alert inventories, clinician collaboration, and structured optimization can create explicit ownership of decision-support problems [23]. More recent stakeholder research on medication-safety recommendations similarly identified clear responsibilities, information exchange, workable protocols, and local collaboration as important implementation conditions [24]. These findings support governance requirements, but not one universal organizational design.

The proposed model therefore treats escalation as a state transition with an accountable destination. **Table 3** makes explicit where ownership should reside as uncertainty narrows and actionability increases.

Table 3. Who owns the signal as medication-safety risk becomes actionable

Risk state	Current owner	Required judgment	Communication route	Time expectation	Failure if ownership is ambiguous
Signal detected; context incomplete	Automated surveillance process	Whether required context has accumulated	Noninterruptive queue or monitored record state	Reassess as relevant data change	Signal persists without review or closure
Clinically plausible risk; patient relevance uncertain	Pharmacist or other designated medication-safety reviewer	Relevance, severity, actionability, and information sufficiency	Documented review with route to responsible clinical team when needed	Priority matched to potential harm and immediacy	Review is performed but no one owns the next action
High-risk, actionable state	Responsible prescriber or clinical team	Treatment change, monitoring, override justification, or immediate mitigation	Direct clinical escalation with acknowledgement	Immediate when delay itself could increase harm	Alert is acknowledged without accountable clinical action
Action initiated but risk unresolved	Named follow-up owner	Whether the risk has resolved, persists, or requires renewed escalation	Closed-loop handoff and documented outcome	Continue until resolution or justified acceptance	Responsibility diffuses across teams or transitions

Note: The escalation states are proposed. The table specifies responsibility logic, not validated staffing models, response-time thresholds, or professional exclusivity.

The proposed risk-adaptive medication-safety model

Medication-alert optimization is already multidimensional: hospitals have used filtering, customization, governance, and other post-implementation strategies rather than one universally effective method [25]. The synthesis proposed here reorganizes those practices around a different question: what response mode is justified by the current preventable-harm risk and uncertainty of a medication-safety signal?

Six dimensions determine that proposed assignment: potential severity, immediacy, preventability, contextual certainty, persistence over time, and the amount of human interpretation required before action. Contextualized drug–drug interaction management has shown greater clinical utility than basic rule-only management in hospitalized patients, supporting the premise that the same nominal interaction can become more or less actionable as patient and treatment context changes [26]. Likewise, development and validation of a dedicated decision-support system for anticoagulant duplication shows that narrowly defined high-risk medication problems can support specific validated rule logic [27]. Real-world adverse-drug-event data can also identify high-risk iatrogenic situations that may inform candidate rules, although such derivation does not establish prospective routing effectiveness [28].

These dimensions are intentionally not combined into a numerical score. A single signal can occupy different states because the dimensions may move in opposite directions. A theoretically severe interaction may remain under surveillance when exposure is remote and context is incomplete; the same interaction may warrant

pharmacist interpretation when patient susceptibility becomes uncertain but clinically important, interruption when an avoidable high-risk exposure is imminent, and direct escalation when exposure has occurred and mitigation cannot wait for routine review. The model therefore classifies the **current decision state**, not the drug pair or rule name. This is also why downgrading is legitimate: when new data reduce severity, resolve uncertainty, or establish a safe monitoring plan, maintaining maximal interruption may add burden without adding equivalent safety value.

The model therefore has four response modes rather than four fixed classes of alert. Silent surveillance retains lower-immediacy or context-dependent signals while new information accumulates. Interruption is reserved for states in which severity, timing, and actionability justify workflow disruption. Pharmacist review is used when automated detection cannot determine patient relevance with sufficient confidence. Clinical escalation transfers an actionable problem to an actor who can mitigate it and close the loop. A signal may move upward or downward among these modes as context changes; no medication, interaction pair, or alert type is permanently assigned to one channel.

Figure 1 depicts this proposed state-dependent migration and its evidence boundaries.

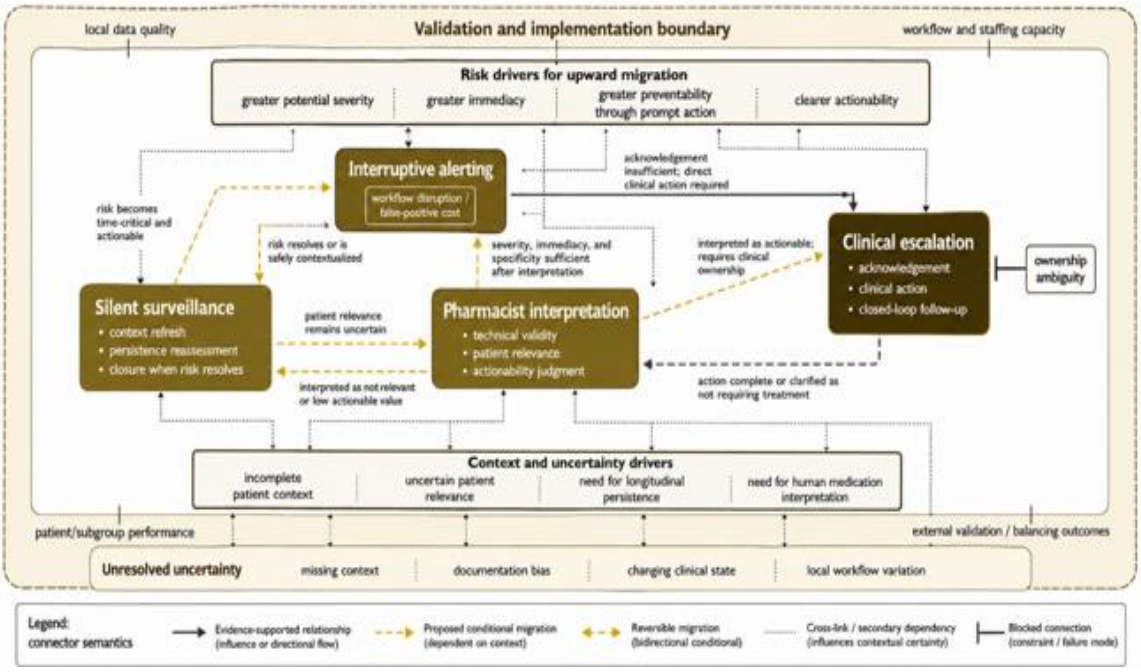

Figure 1. Medication-safety signals migrate among surveillance, interruption, pharmacist interpretation, and escalation as risk and uncertainty change

Calibration, override learning, and alert migration

A risk-adaptive system cannot be calibrated once and considered finished. Alert handling varies with characteristics of the signal, medication, patient, prescriber, and setting, and a recent systematic review and meta-analysis found substantial heterogeneity across those determinants [29]. This argues against a universal override threshold or a single system-wide rule for deciding which alerts deserve interruption. Calibration must instead remain local enough to detect when an apparently useful rule behaves differently across populations, specialties, workflows, or phases of care.

Overrides are particularly easy to misinterpret. A binary override rate cannot distinguish clinically appropriate rejection from redundancy, poor timing, missing context, or genuine habituation. Free-text override comments can provide richer information, and large-language-model methods have been used to summarize those comments for decision-support optimization [30]. Such methods should be treated as tools for pattern detection, not autonomous arbiters of alert suppression or escalation; comments are selectively documented, and model summaries can misrepresent clinician reasoning.

Post-implementation evidence also supports caution. A scoping review of medication-alert optimization found that many interventions reduced alert quantity or improved acceptance and perceived utility, whereas patient-safety outcomes were much less frequently reported [31]. The proposed migration rule is therefore deliberately

reversible. A surveillance signal should move toward interruption only when current evidence justifies increased urgency and actionability; an interruptive alert should move toward surveillance or pharmacist review when repeated local evaluation shows low immediate value, contextual ambiguity, or excessive false-positive cost. Every migration should be accompanied by balancing measures capable of detecting under-alerting, delayed action, or unintended workarounds.

Operationally, calibration requires channel-specific questions. For silent surveillance, the important measures include whether signals are reassessed when relevant data change, whether persistent risk is eventually reviewed, and whether closure is documented. For interruptive alerts, local evaluation should examine not only acceptance but clinical appropriateness of interruption, false-positive cost, workarounds, and missed high-risk states. For pharmacist review, queue size, review timeliness, clinical relevance, and downstream ownership matter. For escalation, acknowledgement without action should not count as successful closure. These measures are proposed evaluation targets rather than established performance standards, and their relative importance will vary with the clinical environment.

Calibration should consequently be understood as governance rather than tuning alone. Local review should combine alert performance, clinical rationale, override explanations, patient-level outcomes where feasible, workflow effects, and evidence of missed risk. The purpose is not to maximize acceptance. It is to maintain an alert channel that remains proportionate to the risk state it is intended to manage.

Failure modes, equity, and boundary conditions

The most obvious failure mode is over-alerting, but a risk-adaptive architecture can fail in the opposite direction. Because alert-fatigue measures are imperfect proxies, reducing visible interruptions does not prove that attention has been better allocated [1]. Silent signals can persist, disappear, or change as medications and clinical data evolve, and those trajectories do not by themselves reveal whether the system caused an appropriate response [15]. Excessive suppression could therefore convert an annoying alert into an invisible unresolved risk. The model requires surveillance states to have explicit reassessment and closure logic.

Patient heterogeneity creates a second boundary. Evidence in older people shows that medication-safety decision support can have different evidential strength for prescribing processes and patient outcomes [3]. Alert handling also varies across patient-, prescriber-, medication-, alert-, and setting-level determinants [29]. These findings do not establish that the proposed model is equitable. They instead require prospective subgroup assessment: a routing strategy should be examined for differential sensitivity, missed-risk patterns, response burden, and access to human review across clinically relevant populations.

A third failure mode is assuming that pharmacist review or escalation capacity exists whenever a rule calls for it. Implementation research identifies responsibility specification, information exchange, protocols, and local collaboration as practical determinants of medication-safety work [24]. A randomized pharmacist intervention can improve selected outcomes in a particular medication-related emergency-care context [21], but that does not make the same service universally available, scalable, or substitutable for local clinical infrastructure. Where pharmacist capacity is constrained, another qualified review pathway must be explicit rather than silently omitted.

Data completeness is another boundary because apparent contextual certainty can be false. Medication lists may omit nonprescription exposure, laboratory values may be stale, indications may be implicit, and information may be fragmented across care settings. A risk-adaptive system that suppresses interruption because a required modifier is missing could therefore create a different failure mode from one that interrupts indiscriminately. Missingness should be treated as an explicit state that can increase uncertainty or trigger review rather than automatically being interpreted as absence of risk.

Equity assessment must similarly extend beyond algorithmic performance. Access to pharmacist interpretation, continuity of records, language-supported communication, and the capacity of clinical teams to respond to escalation can vary across services and settings. The proposed architecture could redistribute burden away from prescribers while unintentionally moving it toward under-resourced review teams or leaving some patients with less complete contextualization. These are plausible implementation risks, not demonstrated effects of the model, and they require prospective measurement rather than assumption.

Finally, improved alert metrics can coexist with uncertain patient benefit. Context-insensitive DDI alerting may fail to reduce interaction-related harm [12], and post-implementation optimization studies still report patient-safety outcomes less often than alert-volume or acceptance outcomes [31]. The proposed architecture should therefore be tested prospectively against preventable medication harm and balancing outcomes, not merely alert

counts. Multisite pragmatic or cluster-randomized evaluations, external validation, subgroup analyses, mixed-method workflow assessment, and prespecified monitoring for missed signals would be appropriate ways to challenge, refine, or reject the model. Until such testing exists, its claims should remain those of a structured Current Opinion: a decision architecture for evaluation, not a validated clinical intervention.

Conclusion

Medication-safety systems should not be judged primarily by how many alerts they generate or eliminate. Alert burden matters because attention is finite, but its clinical significance depends on which signals are interrupted, which remain under surveillance, which require interpretation, and who becomes responsible when a risk becomes actionable.

This Current Opinion proposes a risk-adaptive architecture in which response modality changes with potential severity, immediacy, preventability, contextual certainty, persistence, and interpretive need. Interruption is reserved for circumstances in which workflow disruption is justified by a sufficiently consequential and actionable state. Silent surveillance preserves lower-immediacy or context-dependent signals while new information accumulates. Pharmacist review provides a selective interpretive layer when technical rule firing does not settle patient relevance. Clinical escalation begins when responsibility transfers to a person or team able to act and continues until the risk is resolved or deliberately accepted.

The proposal has important limits. It does not provide validated thresholds, a universal risk score, a proven migration algorithm, or evidence that one professional group should own every ambiguous signal. Its performance would depend on data quality, clinical workflow, staffing, patient population, system configuration, and the availability of closed-loop follow-up. Any implementation should therefore be locally calibrated, reversible, and prospectively evaluated using preventable medication harm, missed-risk events, workflow burden, subgroup performance, and other balancing outcomes.

Most importantly, the model should remain falsifiable. If prospective evaluation shows that a quieter system misses consequential hazards, that pharmacist routing delays action, that escalation ownership is unreliable, or that subgroup performance worsens, the routing logic should be revised or rejected rather than defended by improved alert statistics.

The central shift is conceptual but consequential: medication-safety decision support should optimize the pathway from signal to accountable action, not the count of alerts alone.

Acknowledgments: None

Conflict of Interest: None

Financial Support: None

Ethics Statement: None

References

1. Ray CE, Wilson GM, Hughes AM, Cunningham Goedken C, Liu EPF, Fitzpatrick MA, et al. Alert fatigue measurement in clinical decision support: A systematic review. J Am Med Inform Assoc. 2026;33(8):1523-31. doi:10.1093/jamia/ocag064
2. Graafsma J, Murphy RM, van de Garde EMW, Karapinar-Çarkit F, Derijks HJ, Hoge RHL, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411-22. doi:10.1093/jamia/ocae076
3. Ng Y, Hsu JTY, Ng NNE, Ong JZE, Hsu JLJ, Sulaimi F, et al. Evaluating the role of clinical decision support systems in medication safety for older people: A systematic review. Age Ageing. 2025;54(7):afaf206. doi:10.1093/ageing/afaf206
4. Li L, Baker J, Quirk R, Deidun D, Moran M, Abo Salem A, et al. Drug–drug interactions and actual harm to hospitalized patients: A multicentre study examining the prevalence pre- and post-electronic medication system implementation. Drug Saf. 2024;47(6):557-69. doi:10.1007/s40264-024-01412-w

5. Co Z, Holmgren AJ, Classen DC, Newmark L, Seger DL, Danforth M, et al. The tradeoffs between safety and alert fatigue: Data from a national evaluation of hospital medication-related clinical decision support. J Am Med Inform Assoc. 2020;27(8):1252-8. doi:10.1093/jamia/ocaa098
6. Wright A, Aaron S, Seger DL, Samal L, Schiff GD, Bates DW. Reduced effectiveness of interruptive drug-drug interaction alerts after conversion to a commercial electronic health record. J Gen Intern Med. 2018;33(11):1868-76. doi:10.1007/s11606-018-4415-9
7. Holbrook AM, Matos Silva J, Faruque JAY, Deng J, Schneider T, Jaffer A. Effect of electronic drug-drug interaction alerts on patient and clinician outcomes: A systematic review. J Am Med Inform Assoc. 2025;32(10):1617-28. doi:10.1093/jamia/ocaf139
8. Van De Sijpe G, Walgraeve K, Van Laer E, Quintens C, Machiels C, Foulon V, et al. The impact of customized screening intervals on the burden of drug-drug interaction alerts: An interrupted time series analysis. J Med Syst. 2024;48(1):93. doi:10.1007/s10916-024-02113-8
9. Powers EM, Shiffman RN, Melnick ER, Hickner A, Sharifi M. Efficacy and unintended consequences of hard-stop alerts in electronic health record systems: A systematic review. J Am Med Inform Assoc. 2018;25(11):1556-66. doi:10.1093/jamia/ocy112
10. Bakker T, Klopotowska JE, Dongelmans DA, Eslami S, Vermeijden WJ, Hendriks S, et al. The effect of computerised decision support alerts tailored to intensive care on the administration of high-risk drug combinations, and their monitoring: A cluster randomised stepped-wedge trial. Lancet. 2024;403(10425):439-49. doi:10.1016/S0140-6736(23)02465-0
11. Graafsma J, van de Garde EMW, Derijks HJ, Hoge RHL, Klopotowska JE, Karapinar-Carkit F, et al. Efficiency of computerized clinical decision support systems involving anticoagulants: A flashmob study in Dutch hospital pharmacies. Br J Clin Pharmacol. 2025;91(1):157-65. doi:10.1111/bcp.16236
12. Baysari MT, Hilmer SN, Day RO, Van Dort BA, Zheng WY, Quirk R, et al. Effectiveness of computerised alerts to reduce drug–drug interactions and DDI-related harm in hospitalised patients: A quasi-experimental controlled pre–post study. BMJ Qual Saf. 2025;34(12):788-97. doi:10.1136/bmjqs-2024-018243
13. Jennings LK, Ward RC, Pekar E, Szwast E, Sox L, Hying J, et al. The effectiveness of a noninterruptive alert to increase prescription of take-home naloxone in emergency departments. J Am Med Inform Assoc. 2023;30(4):683-91. doi:10.1093/jamia/ocac257
14. Park H, Chae MK, Jeong W, Yu J, Jung W, Chang H, et al. Appropriateness of alerts and physicians' responses with a medication-related clinical decision support system: Retrospective observational study. JMIR Med Inform. 2022;10(10):e40511. doi:10.2196/40511
15. Bittmann JA, Scherkl C, Meid AD, Haefeli WE, Seidling HM. Event analysis for automated estimation of absent and persistent medication alerts: Novel methodology. JMIR Med Inform. 2024;12:e54428. doi:10.2196/54428
16. Sundermann M, Clendon O, McNeill R, Doogue M, Chin PKL. Optimising interruptive clinical decision support alerts for antithrombotic duplicate prescribing in hospital. Int J Med Inform. 2024;186:105418. doi:10.1016/j.ijmedinf.2024.105418
17. Bauer J, Busse M, Koch S, Schmid M, Sommer J, Fromm MF, et al. Clinical–pharmaceutical assessment of medication CDSS alerts: Content appropriateness and patient relevance in clinical practice. Front Pharmacol. 2025;16:1510425. doi:10.3389/fphar.2025.1510425
18. Flückiger L, Zaugg C, Fiumefreddo R. Impact of pharmacist-evaluated clinical decision support system alerts on potentially missing or inappropriately prescribed proton pump inhibitors at hospital discharge: A retrospective cross-sectional study. Int J Clin Pharm. 2024;46(5):1143-51. doi:10.1007/s11096-024-01746-6
19. Cuvelier E, Robert L, Musy E, Rousselière C, Marcilly R, Gautier S, et al. The clinical pharmacist's role in enhancing the relevance of a clinical decision support system. Int J Med Inform. 2021;155:104568. doi:10.1016/j.ijmedinf.2021.104568
20. Potier A, Dufay E, Dony A, Divoux E, Arnoux LA, Boschetti E, et al. Pharmaceutical algorithms set in a real time clinical decision support targeting high-alert medications applied to pharmaceutical analysis. Int J Med Inform. 2022;160:104708. doi:10.1016/j.ijmedinf.2022.104708
21. Villiet M, Laureau M, Perier D, Pinzani V, Giraud I, Lohan L, et al. Emergency department visits for medication-related events with vs without pharmacist intervention: The URGEIM randomized clinical trial. JAMA Intern Med. 2025;185(6):669-78. doi:10.1001/jamainternmed.2025.0640

22. Jošt M, Kerec Kos M, Kos M, Knez L. Effectiveness of pharmacist-led medication reconciliation on medication errors at hospital discharge and healthcare utilization in the next 30 days: A pragmatic clinical trial. Front Pharmacol. 2024;15:1377781. doi:10.3389/fphar.2024.1377781
23. McCoy AB, Russo EM, Johnson KB, Addison B, Patel N, Wanderer JP, et al. Clinician collaboration to improve clinical decision support: The Clickbusters initiative. J Am Med Inform Assoc. 2022;29(6):1050-9. doi:10.1093/jamia/ocac027
24. Oude Lansink MAM, Vervloet M, van Tholen L, Dankers M, Heringa M, van den Bemt BJF, et al. Barriers and facilitators for the implementation of medication safety recommendations: Focus groups with stakeholders. Int J Clin Pharm. 2026;48(4):1607-21. doi:10.1007/s11096-026-02141-z
25. Van Dort BA, Zheng WY, Sundar V, Baysari MT. Optimizing clinical decision support alerts in electronic medical records: A systematic review of reported strategies adopted by hospitals. J Am Med Inform Assoc. 2021;28(1):177-83. doi:10.1093/jamia/ocaa279
26. Wasylewicz ATM, van de Burgt BWM, Manten T, Kerskes M, Compagner WN, Korsten EHM, et al. Contextualized drug-drug interaction management improves clinical utility compared with basic drug-drug interaction management in hospitalized patients. Clin Pharmacol Ther. 2022;112(2):382-90. doi:10.1002/cpt.2624
27. Dahmke H, Cabrera-Diaz F, Heizmann M, Stoop S, Schuetz P, Fiumefreddo R, et al. Development and validation of a clinical decision support system to prevent anticoagulant duplications. Int J Med Inform. 2024;187:105446. doi:10.1016/j.ijmedinf.2024.105446
28. Dintilhac A, Lohan L, Laureau M, Perier D, Cestac P, Juillard-Condat B, et al. From bedside observations to clinical decision support system (CDSS) rules: Using real-world adverse drug events (ADEs) data to identify high-risk iatrogenic situations. Int J Med Inform. 2026;219:106552. doi:10.1016/j.ijmedinf.2026.106552
29. Raven K, Zhang J, Vagliano I, Jaspers MWM, Peute LW, Medlock S. Determinants influencing medication-related alert handling in the electronic health record: A systematic review and meta-analysis. Int J Med Inform. 2025;203:106011. doi:10.1016/j.ijmedinf.2025.106011
30. Liu S, McCoy AB, Wright AP, Nelson SD, Huang SS, Ahmad HB, et al. Why do users override alerts? Utilizing large language model to summarize comments and optimize clinical decision support. J Am Med Inform Assoc. 2024;31(6):1388-96. doi:10.1093/jamia/ocae041
31. Ledger TS, Brooke-Cowden K, Coiera E. Post-implementation optimization of medication alerts in hospital computerized provider order entry systems: A scoping review. J Am Med Inform Assoc. 2023;30(12):2064-71. doi:10.1093/jamia/ocad193