Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

Why Artificial Intelligence–Enabled Medication Decision Support Becomes Useful, Ignored, or Routinely Overridden: A Realist Review of Workflow Conditions, Professional Trust, and Organizational Responses

Thomas Berger[1]*, Nadia Rahman[2], Lucas Meyer[3], Rania Khalil[4]

1. ***Department of Clinical Pharmacy and AI Decision Support, University of Bern, Bern, Switzerland.***
2. ***Department of Pharmacy Practice and Professional Trust, University of Dhaka, Dhaka, Bangladesh.***
3. ***Department of Pharmaceutical Care and Health Informatics, University of Amsterdam, Amsterdam, Netherlands.***
4. ***Department of Clinical Pharmacy and Organizational Behavior, Cairo University, Cairo, Egypt.***

ARTICLE INFO

Received:
29 April 2026
Received in revised form:
10 September 2026
Accepted:
15 September 2026
Available online:
28 October 2026

Keywords: Artificial intelligence, Medication decision support, Realist review, Alert override, Professional trust, Clinical workflow

ABSTRACT

Artificial intelligence (AI)–enabled medication decision support can improve the identification and prioritization of prescribing risks, yet recommendations may be used selectively, ignored, or routinely overridden. These responses have different clinical meanings that cannot be determined from acceptance rates alone. This realist review examines how workflow conditions, professional trust, and organizational responses interact to shape medication decision-support use. Peer-reviewed literature published between 2018 and 2026 was identified through targeted scholarly searching and mapped against an initial programme theory. Context–mechanism–outcome reasoning was used to distinguish system characteristics, clinician interpretation, and observable actions. The evidence indicates that algorithmic performance, perceived relevance, interruption burden, professional accountability, and institutional implementation practices are distinct considerations in understanding medication-support behavior. It also identifies important differences between medication-alert studies, simulated human–AI decisions, and real-world implementation evaluations. However, complete record-level screening denominators and included-evidence-unit linkage were not recoverable from the search record, limiting final selection-flow reporting and confidence in a comprehensive synthesis. The explanatory account therefore remains conditional: useful uptake requires a combination of clinically relevant advice, workable opportunities for interpretation, and organizational arrangements that preserve professional judgment. Override and non-use require contextual evaluation because they may represent appropriate clinical reasoning, adaptation to workflow constraints, or failure of implementation.

*This is an **open-access** article distributed under the terms of the Creative Commons Attribution-Non Commercial-Share Alike 4.0 License, which allows others to remix, and build upon the work non commercially.*

To Cite This Article: Berger T, Rahman N, Meyer L, Khalil R. Why Artificial Intelligence–Enabled Medication Decision Support Becomes Useful, Ignored, or Routinely Overridden: A Realist Review of Workflow Conditions, Professional Trust, and Organizational Responses. Pharmacophore. 2026;17(5):63-73. https://doi.org/10.51847/H5FK3aoddQ

Introduction

Medication-related clinical decision support is increasingly expected to identify unsafe prescriptions while preserving clinicians' capacity to make timely, patient-specific decisions. Artificial intelligence offers methods for prioritizing alerts and identifying clinically relevant prescribing patterns, but improved classification does not necessarily translate into safer medication use. A scoping review of AI-based medication-alert optimization documented the developing technical evidence while identifying important implementation and evaluation limitations [1]. A subsequent systematic review of electronic drug–drug interaction alerts found that changes in prescribing behavior did not establish corresponding improvements in patient-important outcomes [2]. Moreover, alert fatigue has been defined and measured inconsistently across clinical decision-support research, complicating comparisons between systems that produce different notification volumes, interruption patterns, and

Corresponding Author: Thomas Berger; Department of Clinical Pharmacy and AI Decision Support, University of Bern, Bern, Switzerland. E-mail: thomas.berger@unibe.ch.

professional responses [3]. These distinctions make the clinical meaning of alert acceptance and override an explanatory problem.

The conditions under which clinicians encounter recommendations are central to that problem. Medication-related decision support may arrive during prescribing, dispensing, verification, or medication review, each involving different responsibilities and opportunities for action. A systematic review of clinician-reported acceptance factors identified perceived usefulness, relevance, ease of use, and efficiency as recurring considerations [4]. Professional identity adds another dimension: decision support may be experienced as reinforcing expertise or diminishing autonomy, depending on professional role, task allocation, and organizational arrangements [5]. Trust in AI-supported decisions likewise encompasses judgments about reliability, transparency, and the relationship between the system's recommendations and professional responsibility [6]. These constructs are analytically distinct, although they may interact during a single medication decision.

A recommendation can also receive substantial attention without producing appropriate reliance. Research on AI-assisted clinical decision-making demonstrates that human interaction with algorithmic advice is itself relevant to effectiveness [7]. Medication-alert research suggests that interaction design and clinical role tailoring may reduce unnecessary alert burden [8]. Nevertheless, fewer interruptions cannot independently establish improved medication safety. An alert withheld from a clinician may have been unnecessary, or it may have contained clinically important information. Conversely, a recommendation that is accepted quickly may reflect confident interpretation or insufficient scrutiny. The outcome of interest must therefore extend beyond whether an alert was displayed, acknowledged, accepted, or overridden.

Override behavior offers an especially informative site for examining these distinctions. Liu and colleagues used large language models to summarize clinician comments associated with alert overrides, illustrating how recorded explanations can inform the identification of alert-system problems [9]. Explainable-AI approaches have also been investigated as resources for optimizing clinical decision support [10]. Yet structured override reasons are not equivalent to verified clinical actions. A multisite investigation found considerable variation in the reasons available to users, including entries indicating an intention to monitor or take precautions [11]. Such documentation records a stated response; it does not establish that monitoring occurred, that the original alert was clinically appropriate, or that the resulting medication decision was safe.

Professional trust cannot be treated as an intervention objective independent of recommendation quality. A trust-oriented account of medical AI emphasizes the need to understand how these systems are incorporated into clinical practice [12]. Experimental research involving safe and unsafe explainable-AI recommendations has shown that clinicians' attention to advice and explanations does not provide a simple guarantee of correct decision-making [13]. Additional research identifies individual variability in clinician trust, reliance, and performance when interacting with explainable systems [14]. The distinction between warranted and unwarranted reliance is particularly important in medication management, where following an incorrect recommendation may be consequential and overriding a correct recommendation may also expose patients to harm.

Clinical reasoning supplies the final distinction needed to define the review question. AI should support the interpretation of clinical evidence and decisions rather than merely secure behavioral compliance with its outputs [15]. A machine-learning system developed to identify prescriptions at elevated risk of medication error illustrates the potential of computational prioritization [16]. Whether this capability becomes clinically useful depends on how the result enters professional review, what information is available for verification, and which actions follow. Technical discrimination, clinician interpretation, and medication-related outcomes consequently represent connected but non-equivalent stages of implementation.

This review asks: Under which combinations of workflow conditions, professional trust, and organizational response does AI-enabled medication decision support become useful, selectively used, routinely overridden, worked around, or abandoned? Its realist orientation examines how resources and constraints are interpreted by particular actors in particular settings, and how those interpretations contribute to differentiated outcomes. The central analytical task is to distinguish appropriate clinical challenge from habitual bypass while identifying circumstances in which superficially similar implementations generate different responses.

Materials and Methods

The review was organized around a realist explanatory question concerning the implementation and use of AI-enabled medication decision support. Its scope encompassed medication-alert optimization, prescription-risk prioritization, clinician interaction with recommendations, professional trust, and organizational implementation. Context-sensitive implementation evidence was considered relevant because system performance alone cannot explain how recommendations are incorporated into work. A hospital-focused review informed by the nonadoption, abandonment, scale-up, spread, and sustainability framework identifies technological and organizational considerations relevant to clinical decision-support implementation [17]. Evidence on actual CDSS uptake further establishes the need to distinguish reported intention, opportunity for interaction, and observable system use [18].

The evidence-search period was restricted to publications from 2018 through 2026. Targeted searches used combinations of medication, prescribing, clinical decision support, artificial intelligence, machine learning, override, alert fatigue, workflow, trust, professional identity, and implementation. Bibliographic and publisher records were examined to identify relevant empirical evaluations, evidence syntheses, experimental human–AI studies, and methodological contributions. One PubMed-oriented search combined clinical decision support, medication or prescribing, and artificial intelligence or machine learning

with the specified publication-date restriction. The search process produced a bibliography of thirty-seven selected publications. However, complete source-specific result exports were not preserved, and the pooled retrieval set cannot be reconstructed reliably. Accordingly, the number of records identified, removed as duplicates, screened, excluded, and included as independent evidence units is not reported.

The initial programme theory was developed to examine how the same recommendation might produce different outcomes under different implementation conditions. Realist review is appropriate when the question concerns how interventions generate outcomes through context-sensitive processes rather than only whether an intervention has an average effect [19]. In this review, context denotes the circumstances in which a medication recommendation is encountered, including clinical task, patient-specific information, alert presentation, available time, professional responsibility, and organizational arrangements. A mechanism denotes the response generated when a clinician encounters a resource or constraint—for example, recognizing information as actionable, doubting its relevance, or interpreting a recommendation as conflicting with accountable clinical judgment. An outcome denotes a distinguishable response or consequence, including useful uptake, selective use, override, workaround, or abandonment. These terms describe the analytical approach; the causal status of any particular configuration depends on the supporting evidence.

The initial propositions concerned alert relevance and interruption burden, the relationship between explanation and warranted trust, the influence of professional responsibility, and the organizational handling of recorded overrides. They were treated as candidate explanations rather than predetermined findings. Searching and extraction were directed toward both corroborating and contradictory evidence. The review also distinguished direct medication-support evaluations from general clinical-AI experiments and broader implementation reviews. This distinction was necessary because a mechanism suggested by simulated decision behavior may be relevant to medication practice without demonstrating the same effect during routine prescribing.

Evidence appraisal was organized around relevance, richness, and rigour. In realist methodology, relevance concerns a source's contribution to developing or testing programme theory; richness concerns the explanatory detail available; and rigour concerns the trustworthiness of the evidence supporting a particular inference [20]. These judgments are contribution-specific. An experimental study may provide credible information about reliance on unsafe advice but little evidence about sustained organizational implementation. Conversely, an implementation report may illuminate workload and professional response while offering limited evidence about clinical outcomes. A final source-by-source appraisal could not be completed from the available extraction record, and no aggregate quality score or unsupported confidence rating was assigned.

The unit-of-analysis distinction was preserved during evidence mapping. A bibliographic report was not automatically treated as an independent empirical study. Reviews were used to identify relevant constructs, methodological considerations, and potential primary evidence; their constituent studies were not added independently to a review denominator without record-level verification. Where several publications could describe overlapping data, study-level linkage would be necessary before calculating an included-evidence-unit count. The available record does not establish that linkage. Consequently, the thirty-seven bibliography entries represent the selected citation set and cannot be interpreted as the number of independently included implementation evaluations.

The mapped literature encompasses several analytically different contributions. Some sources examine computational identification or filtering of medication risks; others evaluate alert presentation, clinical responses, or organizational implementation. Experimental human–AI studies address reliance and interpretation under controlled conditions, while evidence syntheses supply broader contextual findings. These differences matter because each source type offers a different opportunity to test the programme theory. **Table 1** distinguishes their implementation settings, explanatory contributions, and limits without treating them as interchangeable studies.

Table 1. Characteristics of mapped evidence contributions and medication decision-support settings relevant to programme-theory testing

Evidence contribution	Medication-support setting or task	Principal analytical contribution	Limitation for CMO interpretation
Medication-alert optimization and prescription-risk modeling	Electronic prescribing, prescription screening, and medication review	Characterizes algorithmic prioritization and potential changes in alert exposure	Predictive performance does not establish enacted clinical utility
Drug-interaction and medication-safety evaluations	Hospital prescribing and medication-alert workflows	Examines alert handling, prescribing responses, and selected clinical outcomes	Local implementation and study design constrain causal interpretation
Override-reason and clinician-feedback studies	Drug-interaction alerts and recorded clinician responses	Provides evidence about documented reasoning and possible alert revision	Stated intentions do not verify subsequent medication-related actions
Human–AI decision experiments	Controlled clinical recommendation tasks	Examines trust, explanation, attention, and reliance on advice	Simulated behavior does not establish sustained routine implementation
Professional and organizational implementation studies	Multidisciplinary clinical decision-support environments	Identifies role, workflow, and institutional conditions relevant to use	Findings from general CDSS require medication-specific qualification

Methodological and secondary evidence syntheses	Cross-setting evidence interpretation	Supports initial theory development, appraisal, and identification of contradictory evidence	Constituent studies require linkage before independent inclusion

Table note: Rows distinguish types of evidence contribution, not mutually exclusive studies or verified included-evidence-unit counts. CMO = context–mechanism–outcome; CDSS = clinical decision support system.

The analysis then mapped source-supported observations to potential context–mechanism–outcome propositions. Descriptive evidence concerning exposure, system interaction, or reported attitudes was distinguished from inferred reasoning mechanisms. In particular, a documented override was not assigned a clinical-appropriateness classification without evidence concerning the recommendation, patient circumstances, stated justification, and subsequent action. The same principle applied to acceptance: following AI advice did not constitute proof of a correct medication decision. Contradictory findings were retained as potential indications of contextual dependence, different outcome definitions, or alternative explanations.

Context–Mechanism–Outcome Synthesis

The mapped literature indicates that AI-enabled medication decision support is shaped first by how a recommendation enters work. Clinicians do not encounter medication advice in an abstract informational environment; they encounter it while prescribing, verifying, dispensing, reviewing, or monitoring treatment. In those settings, the same recommendation may be experienced as timely assistance, an interruptive burden, or an unnecessary challenge to expert judgment. A recent systematic review and meta-analysis of medication-related alert handling identified heterogeneous determinants of response, including alert characteristics, workflow circumstances, and user factors [21]. A prospective study of diagnostic recommendations embedded in medication alerts further suggests that clinicians' responses depend not only on the content of the recommendation but also on whether the alert fits an actionable moment in care [22]. A system-wide observational evaluation of alerts designed to prevent drug-induced long QT syndrome similarly illustrates that the practical meaning of alert exposure depends on clinical setting, task relevance, and the surrounding safety workflow [23].

A second pattern concerns temporal instability. The literature does not support the assumption that acceptance or use is fixed once an AI-enabled support tool has been introduced. A systematic review of clinicians' acceptance and use of decision support over time showed that early implementation responses do not necessarily predict later routine engagement [24]. Task-specific applications may perform differently from general alert systems because they focus on a narrower, more interpretable problem; an example is machine-learning support aimed at wrong-drug and look-alike/sound-alike medication risks [25]. By contrast, adoption-oriented work with pharmacists indicates that intention to use, perceived performance benefit, and facilitating conditions may be related without being equivalent to observed clinical use [26]. Together, these sources suggest that use must be treated as a situated and evolving outcome rather than a stable property of the technology.

The next pattern concerns organizational response and the clinical meaning of alerts after deployment. In a controlled pre–post study, computerized alerts reduced drug–drug interactions and related harm under a concrete organizational implementation arrangement [27]. That finding is important because it moves beyond simple interaction metrics. However, other evidence indicates that apparently supportive design features can affect clinician behavior in ways that are not straightforwardly beneficial. Comparative work on explanation formats shows that different forms of explanation can alter decisions, but explanation quality and explanation fit are separate issues [28]. Large language model work in medication safety also demonstrates strong capability within bounded tasks, yet task performance in controlled environments does not by itself establish routine effectiveness at the point of care [29]. Organizational context therefore affects not only whether a system is used, but how users interpret the legitimacy, urgency, and trustworthiness of its advice.

A fourth pattern broadens the context beyond clinicians alone. Older adults' perceptions of AI tools for medication decisions indicate that patient-facing trust and acceptability can become relevant boundary conditions where medication decisions are discussed with patients or caregivers [30]. A scoping review of generative AI and large language models in medication-related harm demonstrates rapid expansion of proposed applications but variable evidentiary maturity across those applications [31]. A deployment framework oriented to real-world clinical implementation reinforces the need to separate model development from successful organizational embedding [32]. These sources do not imply that patient perceptions and implementation science should be conflated. Rather, they indicate that a medication recommendation may pass through several interpretive environments before it becomes action.

The burden created by repeated alerts remains a central contextual condition. A qualitative study of alert fatigue found that clinicians describe interruptions, repetition, and difficulty distinguishing important from low-value notifications as features of their everyday work environment [33]. Importantly, infrastructure changes can alter these responses even when the nominal decision-support purpose remains the same. After conversion to a commercial electronic health record, the effectiveness of interruptive drug–drug interaction alerts declined, illustrating that workflow architecture and platform design can change how recommendations are handled [34]. These findings support a contextual interpretation of override behavior: what appears to be resistance to support may instead be an adaptation to a system whose timing, salience, or integration into work has changed.

The most revealing evidence comes from studies that distinguish technical prioritization from clinical interpretation. A probabilistic machine-learning medication support system reduced prescribing errors and adverse drug events in one inpatient implementation, but that result remained bound to a specific design and local workflow [35]. Research on renal medication alerts showed that override rates alone do not explain whether the overrides were appropriate, because the significance of an override depends on clinical context and the structure of the alert itself [36]. Work on machine-learning approaches to filtering

medication alerts indicates that reducing alert burden may be technically feasible, but whether filtered alerts preserve safe decision-making remains an implementation question rather than a purely computational one [37]. These findings help differentiate at least three distinct outcomes: useful uptake, in which a recommendation supports action; selective use or selective override, in which the recommendation is challenged but remains part of the decision process; and routine override or abandonment, in which the recommendation ceases to function as a meaningful resource.

What emerges across these sources is not a single trajectory from system exposure to system acceptance. Instead, alert exposure, perceived relevance, explanation, accountability, and organizational follow-up appear to combine in different ways across settings. The analytical task is therefore to identify recurring configurations rather than average effects. The configurations below capture how exposure to medication advice may be shaped before clinicians decide whether to act on it. A clearer view of these configurations is provided by the patterned comparison below, which focuses on the contextual arrangements that shape how AI-enabled medication advice is encountered and interpreted.

Table 2 summarizes context configurations shaping exposure to and interpretation of AI-enabled medication decision support.

Table 2. Context configurations shaping exposure to and interpretation of AI-enabled medication decision support

Context configuration	Exposure conditions	Likely reasoning or interpretive response	Observable behavioral pattern	Key boundary condition
Timely, task-aligned recommendation	Alert or recommendation appears during a clinically actionable step and is clearly linked to the immediate medication task	Advice is recognized as relevant and worth checking	Considered use or prompt verification	Requires sufficient information to judge the recommendation
High-volume interruptive alert environment	Repeated alerts compete with other tasks and interrupt workflow	Advice is discounted as routine noise or low-value interruption	Rapid dismissal, habitual override, or minimal engagement	Frequency alone does not establish inappropriateness of the alert
Narrowly targeted safety application	Decision support addresses a specific, recognizable medication risk	Advice appears more intelligible and clinically bounded	Higher likelihood of selective uptake	Transferability to broader prescribing contexts may be limited
Explanation-rich but quality-variable recommendation	Recommendation is accompanied by explanation or rationale, but correctness remains uncertain	Explanation may increase confidence without guaranteeing warranted trust	Use, over-reliance, or selective challenge	Explanation fit and recommendation validity must be distinguished
Role-misaligned recommendation	Support is delivered in a way that does not fit the clinician's responsibility or task authority	Advice is interpreted as administratively intrusive or operationally irrelevant	Override, deferral, or workaround	Professional role and accountability differ across settings
Workflow under strain or staffing pressure	Time, cognitive load, or competing priorities limit reflective engagement	Resource-preserving reasoning dominates and recommendations receive abbreviated appraisal	Bypass or selective attention	Behavior may reflect workload conditions rather than attitude toward AI
Organization responsive to user feedback	Recorded comments, override patterns, or safety data are reviewed and used to revise rules or presentation	Advice becomes more credible or easier to interpret over time	Shifts from routine override toward selective use	Requires institutional capacity for revision and learning
Infrastructure disruption or platform change	EHR migration or interface redesign alters the way alerts are surfaced	Previously familiar support is reinterpreted under new usability conditions	Changed override patterns or degraded use	Apparent decline may stem from system change, not loss of trust

The mapped evidence also clarifies that mechanisms should not be inferred directly from observable actions. An override can result from informed disagreement, time pressure, poor recommendation fit, or routine habituation. Likewise, a click consistent with acceptance can represent careful professional judgment or uncritical reliance. To understand these differences, the synthesis must consider how clinicians interpret the informational and organizational resources embedded in the recommendation.

First, recommendations appear more likely to become useful when clinicians can relate them to the immediate clinical task and can assess them without excessive disruption. The relevant mechanism is not "acceptance" in itself but recognition that the recommendation is actionable and worth cognitive effort. Studies of alert determinants, embedded diagnostic advice, and role-specific systems all support the importance of this contextual fit [21, 22, 25].

Second, professional trust operates less as a stable attitude than as a calibration process. Explanation, prior experience, and recommendation quality influence whether the clinician regards the advice as a serious aid, a prompt requiring verification, or an insufficiently reliable system output. Experimental and conceptual work on explanation and trust supports this interpretation [28, 29]. That calibration can fail in both directions: the clinician may discount a useful recommendation because previous alerts have been noisy, or may give unwarranted weight to a persuasive but inaccurate recommendation.

Third, organizational response influences whether local non-use is an endpoint or a source of learning. Where override comments, observed patterns, and adverse outcomes feed back into rule revision, interface redesign, or workflow adaptation, the organization treats clinician behavior as diagnostic information rather than simple noncompliance. Evidence from implementation and alert-optimization studies suggests that this feedback capacity is central to sustained usefulness [22, 27]. Where such organizational learning is absent, recurrent override may harden into an enduring pattern of disregard.

These observations support three broad CMO propositions. In contexts where medication advice is timely, task-aligned, and supported by interpretable information, clinicians may reason that engagement is worthwhile, producing selective or useful uptake. In contexts where advice arrives as a repetitive interruption with uncertain relevance, clinicians may reason that attention is better conserved, producing routine override or workaround. In contexts where organizations revise systems in response to clinician feedback, users may come to treat recommendations as part of an adaptive safety process rather than an external imposition, increasing the likelihood of calibrated engagement. The next section refines these propositions into a programme theory that more clearly distinguishes use, selective override, routine override, workaround, and abandonment.

Refined Programme Theory

The synthesis supports a more differentiated account than the initial assumption that AI-enabled medication support succeeds when it is accurate and fails when it is overridden. Technical accuracy remains important, but the evidence suggests that medication decision support becomes useful only when computational performance, workflow opportunity, professional interpretation, and organizational learning align. A recommendation is therefore not simply "adopted" or "rejected"; it is processed within a sociotechnical setting that shapes what the recommendation means and what the clinician can realistically do with it.

The refined programme theory begins with workflow conditions. These include timing within the medication task, interaction burden, role alignment, information availability, and competing operational demands. Workflow conditions determine whether the recommendation presents itself as a usable resource or as an interruption. When the alert or recommendation fits the clinical moment and can be interpreted without excessive friction, it creates the possibility of meaningful engagement. When it arrives at a poorly aligned moment or imposes repeated cognitive cost, it makes disregard or rapid dismissal more likely.

These workflow conditions activate trust-related mechanisms, but not in the sense of generalized positive attitude. Trust in this review is better understood as a calibrated judgment about whether the recommendation is sufficiently relevant, intelligible, and accountable to influence action. Explanations, previous experience with the system, the apparent seriousness of the risk, and the clinician's own sense of responsibility all contribute to that judgment. Calibrated trust supports selective use: the clinician takes the recommendation seriously enough to inspect or test it, but not so uncritically as to surrender professional judgment. Miscalibrated trust produces two opposite outcomes. Under-trust may cause useful recommendations to be ignored because the system is perceived as noisy or low-value. Over-trust may lead clinicians to give excessive weight to persuasive or well-explained recommendations that are nevertheless incorrect.

Organizational response then determines whether these local interactions accumulate into improvement or deterioration. Organizations that review overrides, learn from comments, adjust rules, and redesign workflows can convert clinician responses into system refinement. In such settings, override is not merely treated as failure; it can signal where the recommendation lacks relevance or where the workflow imposes avoidable burden. Organizations that do not perform this learning function allow misalignment to persist. In those circumstances, routine override can become normalized, and even potentially useful recommendations may lose their practical authority.

The resulting outcome categories are analytically distinct. Useful uptake occurs when recommendations fit the task, are regarded as worth checking, and can be incorporated into action or verification. Selective use occurs when the recommendation enters reasoning but is either qualified or challenged in light of patient specifics, alternative evidence, or contextual constraints. Routine override occurs when clinicians encounter the recommendation as a familiar interruption and no longer treat it as a meaningful informational resource. Workaround occurs when users adapt their behavior to manage the system rather than the clinical problem, for example by dismissing or bypassing recommendations to preserve workflow continuity. Abandonment occurs when the system or recommendation class loses sustained relevance in practice. These categories should not be collapsed into a single acceptance continuum.

The theory also clarifies how apparently similar implementations can produce different outcomes. Two organizations may deploy comparable recommendation engines, yet one experiences useful selective use while the other experiences rapid override. The difference may lie not in the algorithm alone, but in whether recommendations are delivered at an actionable point, whether users can inspect the underlying rationale, whether alert burden is already high, and whether feedback triggers rule revision. The key explanatory move is thus configurational: outcomes depend on the relationship among workflow conditions, trust calibration, and organizational response, rather than on any one element in isolation.

The refined patterns can be summarized as distinct mechanism–outcome combinations that preserve the differences among useful uptake, selective use, override, workaround, and abandonment.

Table 3 summarizes mechanism–outcome patterns distinguishing useful uptake, selective use, routine override, workaround, and abandonment.

Table 3. Mechanism–outcome patterns distinguishing useful uptake, selective use, routine override, workaround, and abandonment

Outcome pattern	Typical contextual starting point	Dominant reasoning or mechanism	Observable decision response	Programme-theory implication
Useful uptake	Recommendation is timely, clinically relevant, and interpretable	Clinician judges the advice to be actionable and worth integrating into care	Advice informs or changes action in a purposeful way	Usefulness depends on fit and interpretability, not mere availability
Selective use	Recommendation is relevant but requires contextual qualification	Clinician treats the advice as an input to judgment rather than a command	Recommendation is checked, partially used, or conditionally overridden	Appropriate use may include challenge and qualification
Routine override	Recommendation is repetitive, weakly relevant, or poorly timed	Clinician conserves attention by discounting the alert as low-value	Rapid dismissal or repeated override	Persistent misalignment erodes practical authority of the system
Workaround	System burden interferes with completing clinical work	Clinician manages the workflow burden rather than the recommendation content	Bypass, deferral, or tactics to reduce system friction	Behavior reflects system adaptation pressure more than evaluation of recommendation quality
Abandonment	Support remains insufficiently integrated and unrevised over time	Clinician or organization concludes the system offers limited practical value	Sustained non-use or effective disappearance from routine work	Organizational non-learning allows local problems to become systemic non-use
Unsafe over-reliance	Recommendation appears authoritative or well explained despite possible error	Clinician grants excessive credibility to advice	Advice is followed without adequate checking	"Use" is not necessarily beneficial and requires quality-sensitive interpretation

The final programme theory is shown in **Figure 1**. It preserves the article's central claim that medication decision support becomes useful, ignored, or overridden not because one outcome is the opposite of another, but because different contexts trigger different interpretive processes and organizational reactions.

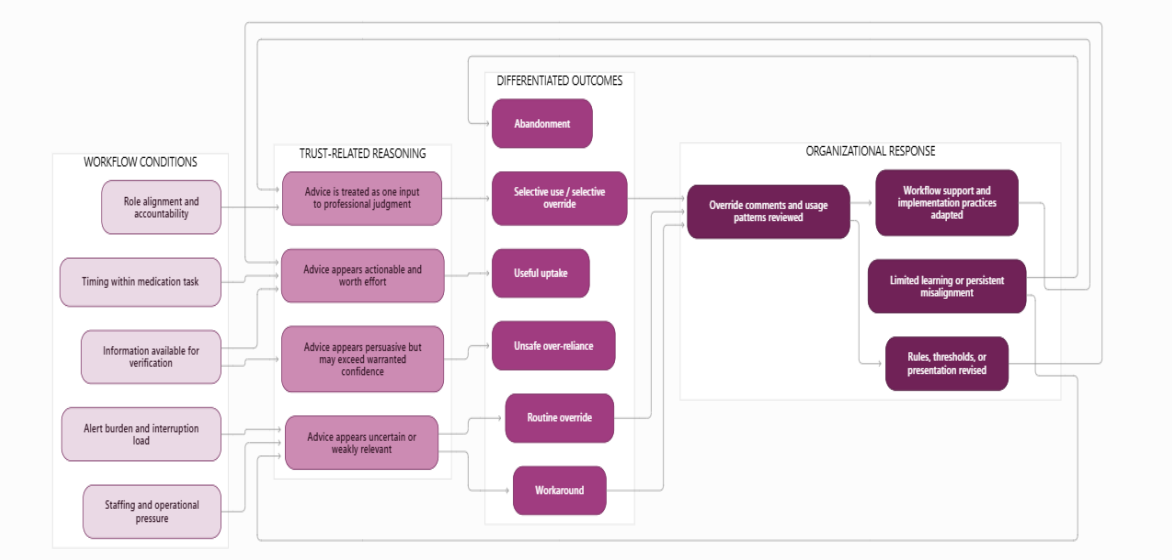


Figure 1. How workflow fit, trust calibration, and organizational learning sort medication advice into use, challenge, or disregard

Results and Discussion

This review addresses a problem often obscured in discussions of medication decision support: observable use, override, and non-use do not speak for themselves. The mapped evidence suggests that AI-enabled medication recommendations acquire clinical meaning only within concrete workflows, under conditions of professional accountability, and in relation to organizational learning. This shifts the explanatory focus away from generic adoption language and toward the circumstances in which clinicians find medication advice worth acting on, worth challenging, or no longer worth attending to.

The principal contribution of the review is therefore conceptual and organizational rather than purely technical. Existing evidence already indicates that alert quality, alert burden, and implementation circumstances matter [1, 2, 21]. What the realist synthesis adds is an account of how these factors combine. Workflow fit influences whether the recommendation can enter

reasoning as a usable resource. Trust calibration influences whether the clinician grants the advice appropriate weight. Organizational response determines whether patterns of override and feedback lead to refinement or to entrenched disregard. This combination explains why routine override cannot automatically be treated as resistance and why high use cannot automatically be treated as success.

A second contribution is the distinction between appropriate selective override and habitual bypass. This distinction has practical significance for both evaluation and governance. Studies of override comments and coded reasons show that clinicians often document explanations for dismissing or qualifying alerts [9, 11]. Yet those explanations do not verify the appropriateness of the final medication decision. More informative evaluation would therefore connect the recommendation, the documented rationale, subsequent clinical action, and where possible the patient-level consequence. Without that chain, both adoption-oriented and safety-oriented interpretations remain partial. Controlled implementation studies and setting-specific medication alert evaluations support the importance of going beyond simple click-based measures [23, 27, 36].

The synthesis also suggests that explanation should be treated cautiously. Explainable outputs may improve interpretability and can support alert optimization or more comprehensible recommendations [10, 28]. However, experimental human–AI evidence indicates that explanation does not automatically secure warranted reliance [13, 14]. An explanation that is clear but attached to a flawed recommendation can increase the persuasive force of error rather than reduce it. For medication decision support, the important outcome is not the presence of explanation, but whether explanation helps clinicians perform accountable checking under real workflow conditions.

Organizational learning emerged as a particularly important feature of sustained usefulness. Systems that remain static in the face of recurring override, workflow burden, or shifting infrastructure are more likely to become normalized as noise. By contrast, organizations that review override patterns, examine comments, and revise rules or presentation are better positioned to turn local non-use into information for redesign [9, 22]. This point matters because medication decision support is not implemented into a static environment. Staffing pressures, changing formularies, new EHR configurations, and broader clinical priorities can alter the experience of a recommendation even when the underlying decision logic remains unchanged [33, 34].

These observations have several implications for implementation and evaluation. First, studies should separate technical performance from clinical usefulness. Model discrimination or alert filtering performance is relevant, but it is not a sufficient endpoint for medication decision support [16, 37]. Second, evaluations should preserve the difference between use and appropriate use. The most useful process metrics will capture not only exposure and response, but also downstream verification and action. Third, implementation teams should treat override patterns as signals requiring interpretation rather than as simple indicators of poor compliance. Fourth, explanation design should be tested alongside workload, timing, and organizational revision processes rather than in isolation [24, 32].

The review also has substantial limitations. The evidence base is heterogeneous in design, outcome definition, and clinical context. The mapped literature includes secondary reviews, simulation-based human–AI studies, observational implementations, and a smaller number of more direct medication decision-support evaluations. Complete record-level search exports were not preserved, so full retrieval denominators, screening counts, and a definitive included-evidence-unit total could not be recovered. Contribution-level appraisal by relevance, richness, and rigour was therefore only partially completed. This limits the certainty with which any one CMO configuration can be advanced as a settled explanation. The programme theory should accordingly be interpreted as a refined and evidence-bounded account rather than a definitive causal map [19, 20].

Transferability must also be qualified. Findings drawn from medication alerts, pharmacist-oriented applications, patient-facing surveys, and broader AI decision studies should not be treated as interchangeable. The strongest transfer appears where the core sociotechnical problem is shared: a recommendation enters practice, must be interpreted under accountability, and is embedded within an organization capable of either learning or not learning from use patterns. Even there, professional role, care setting, alert type, and platform design can meaningfully alter the response [17, 18, 26].

The boundary conditions of the review can be summarized as configurations that support or weaken the refined theory's applicability.

Table 4 summarizes refined context–mechanism–outcome configurations with programme-theory implications and boundary conditions.

Table 4. Refined context–mechanism–outcome configurations with programme-theory implications and boundary conditions

Refined configuration	Programme-theory implication	Likely outcome tendency	Boundary condition limiting transfer
Timely recommendation + interpretable rationale + actionable workflow moment	Medication advice can function as a practical cognitive resource	Useful uptake	Requires adequate information and manageable interruption load
Relevant recommendation + strong professional accountability + opportunity for checking	Clinical challenge becomes part of appropriate use	Selective use or selective override	Requires space for judgment rather than mandatory compliance logic
Repetitive low-value alerting + high workload + poor fit with task	Attention-conserving reasoning dominates	Routine override	Similar override rates may mask different clinical appropriateness

Persuasive recommendation + explanation that appears credible + limited checking opportunity	Confidence can exceed warranted trust	Unsafe over-reliance	Depends on recommendation quality and local safeguards
Override feedback reviewed + organizational willingness to revise	Non-use can become a learning signal	Movement toward more selective, calibrated use	Requires institutional capacity for monitoring and redesign
Infrastructure change + unchanged rule burden + limited adaptation	Established alert behavior can deteriorate despite nominally similar support	Workaround or abandonment	Platform effects may dominate perceptions of AI quality
Narrowly targeted medication safety task + bounded risk domain	Task-specific design may preserve intelligibility and relevance	More stable selective uptake	Transfer to broad prescribing or multi-alert environments is uncertain
Adoption measured only through exposure or clicks	Evaluation underestimates differences between challenge, use, and disregard	Misleading interpretation of success or failure	Requires linkage to action and clinical context for meaningful inference

Future work should therefore test these configurations more directly. Stronger studies would link the recommendation to the clinical task, capture the user's stated reasoning, examine downstream action, and assess whether organizational feedback changes the system over time. They would also compare settings in which similar medication recommendations produce different override and use patterns, thereby clarifying whether the critical difference lies in alert logic, explanation, workload, or implementation practice. Such studies would help determine which parts of the present programme theory are robust and which remain provisional.

Conclusion

AI-enabled medication decision support becomes useful, ignored, or routinely overridden through context-sensitive interactions among workflow fit, professional trust calibration, and organizational response. Medication advice is most likely to support care when it reaches clinicians at an actionable moment, can be interpreted without excessive burden, and is embedded in organizations that learn from patterns of use and override. Non-use and override are not inherently failures, because they may represent appropriate clinical challenge, adaptation to burdensome workflow conditions, or the consequence of poorly aligned implementation. The strongest evidence-bounded conclusion is therefore conditional: the practical value of AI-supported medication advice depends less on technical capability alone than on whether the surrounding sociotechnical system enables clinicians to treat the advice as a meaningful, accountable, and revisable resource.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. 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
2. 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
3. Ray CE, Wilson GM, Hughes AM, Goedken CC, 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
4. Westerbeek L, Ploegmakers KJ, de Bruijn GJ, Linn AJ, van Weert JCM, Daams JG, et al. Barriers and facilitators influencing medication-related CDSS acceptance according to clinicians: A systematic review. Int J Med Inform. 2021;152:104506. doi:10.1016/j.ijmedinf.2021.104506
5. Ackerhans S, Huynh T, Kaiser C, Schultz C. Exploring the role of professional identity in the implementation of clinical decision support systems—A narrative review. Implement Sci. 2024;19(1):11. doi:10.1186/s13012-024-01339-x
6. Tun HM, Rahman HA, Naing L, Malik OA. Trust in artificial intelligence–based clinical decision support systems among health care workers: Systematic review. J Med Internet Res. 2025;27. doi:10.2196/69678

7. Jackson NJ, Brown KE, Miller R, Murrow M, Cauley MR, Collins BX, et al. Factors influencing the effectiveness of artificial intelligence-assisted decision-making in medicine: A scoping review. J Am Med Inform Assoc. 2026;33(5):1054-64. doi:10.1093/jamia/ocag002
8. Hussain MI, Reynolds TL, Zheng K. Medication safety alert fatigue may be reduced via interaction design and clinical role tailoring: A systematic review. J Am Med Inform Assoc. 2019;26(10):1141-9. doi:10.1093/jamia/ocz095
9. 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
10. Liu S, McCoy AB, Peterson JF, Lasko TA, Sittig DF, Nelson SD, et al. Leveraging explainable artificial intelligence to optimize clinical decision support. J Am Med Inform Assoc. 2024;31(4):968-74. doi:10.1093/jamia/ocae019
11. Wright A, McEvoy DS, Aaron S, McCoy AB, Amato MG, Kim H, et al. Structured override reasons for drug-drug interaction alerts in electronic health records. J Am Med Inform Assoc. 2019;26(10):934-42. doi:10.1093/jamia/ocz033
12. Zuchowski LC, Zuchowski ML, Nagel E. A trust based framework for the envelopment of medical AI. NPJ Digit Med. 2024;7(1):230. doi:10.1038/s41746-024-01224-3
13. Nagendran M, Festor P, Komorowski M, Gordon AC, Faisal AA. Eye tracking insights into physician behaviour with safe and unsafe explainable AI recommendations. NPJ Digit Med. 2024;7(1):202. doi:10.1038/s41746-024-01200-x
14. Nicolson A, Bradburn E, Gal Y, Papageorghiou AT, Noble JA. The human factor in explainable artificial intelligence: Clinician variability in trust, reliance, and performance. NPJ Digit Med. 2025;8(1):658. doi:10.1038/s41746-025-02023-0
15. Sokol K, Fackler J, Vogt JE. Artificial intelligence should genuinely support clinical reasoning and decision making to bridge the translational gap. NPJ Digit Med. 2025;8(1):345. doi:10.1038/s41746-025-01725-9
16. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O, et al. A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. J Am Med Inform Assoc. 2020;27(11):1688-94. doi:10.1093/jamia/ocaa154
17. Abell B, Naicker S, Rodwell D, Donovan T, Tariq A, Baysari M, et al. Identifying barriers and facilitators to successful implementation of computerized clinical decision support systems in hospitals: A NASSS framework-informed scoping review. Implement Sci. 2023;18(1):32. doi:10.1186/s13012-023-01287-y
18. Kouri A, Yamada J, Cheung JLS, Van de Velde S, Gupta S. Do providers use computerized clinical decision support systems? A systematic review and meta-regression of clinical decision support uptake. Implement Sci. 2022;17(1):21. doi:10.1186/s13012-022-01199-3
19. Duddy C, Wong G. Grand rounds in methodology: When are realist reviews useful, and what does a 'good' realist review look like? BMJ Qual Saf. 2023;32(3):173-80. doi:10.1136/bmjqs-2022-015236
20. Dada S, Dalkin S, Gilmore B, Hunter R, Mukumbang FC. Applying and reporting relevance, richness and rigour in realist evidence appraisals: Advancing key concepts in realist reviews. Res Synth Methods. 2023;14(3):504-14. doi:10.1002/jrsm.1630
21. 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
22. Liu YC, Lin GL, Scholl J, Hung YC, Lin YJ, Li YC, et al. Evaluation of diagnostic recommendations embedded in medication alerts: Prospective single-arm interventional study. J Med Internet Res. 2025;27. doi:10.2196/70731
23. Trinkley KE, Simon ST, Rosenberg MA. Impact of an alert-based inpatient clinical decision support tool to prevent drug-induced long QT syndrome: Large-scale, system-wide observational study. J Med Internet Res. 2025;27. doi:10.2196/68256
24. Newton N, Bamgboje-Ayodele A, Forsyth R, Tariq A, Baysari MT. A systematic review of clinicians' acceptance and use of clinical decision support systems over time. NPJ Digit Med. 2025;8(1):309. doi:10.1038/s41746-025-01662-7
25. Chen CY, Chen YL, Scholl J, Yang HC, Li YCJ. Ability of machine-learning based clinical decision support system to reduce alert fatigue, wrong-drug errors, and alert users about look alike, sound alike medication. Comput Methods Programs Biomed. 2024;243:107869. doi:10.1016/j.cmpb.2023.107869
26. Xie Y, Song Z, Wang E, Li P, Wang J. Beyond willingness: Unpacking pharmacists' adoption of AI-driven clinical decision support systems through an extended UTAUT framework. Front Public Health. 2026;14:1728867. doi:10.3389/fpubh.2026.1728867
27. Baysari MT, Hilmer SN, Day RO, Van Dort BA, Zheng WY, Quirk R, et al. Effectiveness of computerised alerts to reduce drug–drug interactions (DDIs) 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
28. Hur S, Lee Y, Park J, Jeon YJ, Cho JH, Cho D, et al. Comparison of SHAP and clinician friendly explanations reveals effects on clinical decision behaviour. NPJ Digit Med. 2025;8(1):578. doi:10.1038/s41746-025-01958-8
29. Ong JCL, Jin L, Elangovan K, Lim GYS, Lim DYZ, Sng GGR, et al. Large language model as clinical decision support system augments medication safety in 16 clinical specialties. Cell Rep Med. 2025;6(10):102323. doi:10.1016/j.xcrm.2025.102323

30. Vordenberg SE, Nichols J, Marshall VD, Weir KR, Dorsch MP. Investigating older adults' perceptions of AI tools for medication decisions: Vignette-based experimental survey. J Med Internet Res. 2024;26. doi:10.2196/60794
31. Ong JCL, Chen MH, Ng N, Elangovan K, Tan NYT, Jin L, et al. A scoping review on generative AI and large language models in mitigating medication related harm. NPJ Digit Med. 2025;8(1):182. doi:10.1038/s41746-025-01565-7
32. You JG, Hernandez-Boussard T, Pfeffer MA, Landman A, Mishuris RG. Clinical trials informed framework for real world clinical implementation and deployment of artificial intelligence applications. NPJ Digit Med. 2025;8(1):107. doi:10.1038/s41746-025-01506-4
33. Newton N, Bamgboje-Ayodele A, Forsyth R, Tariq A, Huang J, Yannam R, et al. Experiences of alert fatigue and its contributing factors in hospitals: Qualitative study. J Med Internet Res. 2026;28. doi:10.2196/78676
34. 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
35. Segal G, Segev A, Brom A, Lifshitz Y, Wasserstrum Y, Zimlichman E. Reducing drug prescription errors and adverse drug events by application of a probabilistic, machine-learning based clinical decision support system in an inpatient setting. J Am Med Inform Assoc. 2019;26(12):1560-5. doi:10.1093/jamia/ocz135
36. Shah SN, Amato MG, Garlo KG, Seger DL, Bates DW. Renal medication-related clinical decision support (CDS) alerts and overrides in the inpatient setting following implementation of a commercial electronic health record: Implications for designing more effective alerts. J Am Med Inform Assoc. 2021;28(6):1081-7. doi:10.1093/jamia/ocaa222
37. Liu S, Kawamoto K, Del Fiol G, Weir C, Malone DC, Reese TJ, et al. The potential for leveraging machine learning to filter medication alerts. J Am Med Inform Assoc. 2022;29(5):891-9. doi:10.1093/jamia/ocab292