Research Article

Predicting the Probability That Open-Access Clinical Literature Saves Lives

David Cawthorpe[1]

1. Cumming School of Medicine, University of Calgary, Canada

Whether open-access (OA) clinical literature directly saves lives is frequently debated, yet empirical documentation is scarce because clinical notes rarely record how evidence was accessed. This study synthesizes high-impact cases of OA-enabled clinical change—most notably the SARS-CoV-2 PCR diagnostic protocol and the RECOVERY dexamethasone findings—and develops an expanded Bayesian predictive model estimating the probability that a single clinician reading one OA article saves a life. We integrate three primary evidence bases: (1) clinician-reported rates of practice change following article consultation, (2) the proportion of clinical decisions that influence short- or long-term mortality, and (3) empirically observed mortality reductions following OA-mediated dissemination of life-saving therapeutic evidence. We then extend this model by incorporating additional determinants of diagnostic and therapeutic accuracy, including medical error rates, years of clinical experience, multimorbidity-dependent diagnostic entropy, cognitive load, structural barriers, team-based reliability, guideline adherence, and electronic health record (EHR)–related error susceptibility, formalized in a multilevel Bayesian framework. The core model yields a probability range of $p \approx 0.003–0.02$ that a clinician–article encounter prevents one death, corresponding to a Number Needed to Treat (NNT) analog of approximately 50–330 clinician–article encounters. After accounting for heterogeneity in clinical acuity, multimorbidity, and the extended set of clinician and system parameters, hierarchical Bayesian extensions adjust the predictive interval to $p \approx 0.002–0.03$ and NNT $\approx 30–500$. The integrated analysis demonstrates that OA literature meaningfully increases the probability of life-saving clinical decisions, especially in high-acuity environments where marginal improvements in evidence latency and accuracy have large mortality consequences.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

 doi.org/10.32388/U2SYIR

Introduction

Open-access scientific literature is intended to accelerate the translation of evidence into clinical practice, yet it is often claimed that no case exists in which access to a full-text OA paper saved a life. This framing relies on a misunderstanding of evidence diffusion and clinical documentation: clinicians do not typically record whether information was accessed through OA or subscription channels, so access modality is largely invisible in the chart. Nevertheless, empirical examples show that OA dissemination has produced measurable survival benefits at population scale. The SARS-CoV-2 PCR diagnostic protocol, published OA in *Eurosurveillance* in January 2020, enabled laboratories around the world to implement accurate testing within days, a critical early-pandemic intervention[1]. Similarly, the RECOVERY trial's dexamethasone findings, released first as an OA medRxiv preprint, produced immediate global changes in clinical practice that were temporally associated with reductions in mortality[2][3][4]. Evidence from HINARI/Research4Life demonstrates that expanded OA access improves diagnostic accuracy, guideline adherence, and evidence-based treatment in low-resource settings[5][6]. Evaluating whether OA literature "saves lives" therefore requires probabilistic modeling rather than anecdotal case documentation. Given that clinical decision-making is probabilistic, heterogeneous, and distributed across clinicians and settings, the relevant question is: what is the probability that one OA clinical paper read by one clinician contributes to a decision that prevents at least one death?

Methods

This study integrates empirical data and a multilevel Bayesian framework to estimate the probability that one OA article read by one clinician saves a life. The starting point is a three-factor decomposition. First, multiple studies show that approximately 25–45% of clinicians modify clinical decisions after reading a relevant article. Second, epidemiological data indicate that 5–15% of hospital-based clinical decisions directly affect short-term mortality, with up to 30% influencing long-term survival in chronic disease. Third, documented OA-mediated survival effects—such as the RECOVERY dexamethasone and SARS-CoV-2 PCR dissemination events—demonstrate concrete cases where OA accelerated adoption of life-saving interventions. These domains are combined in an initial multiplicative model:

$$p = P\left(\text{practice change}\right) \times P\left(\text{decision affects mortality}\right) \times P\left(\text{correct application of evidence}\right).$$

The model defines p as the probability that a clinician–article encounter prevents a death. We model individual variation in p using a hierarchical structure in which the log-odds of clinician-specific probabilities p_i follow $\text{logit}(p_i) \sim \mathcal{N}(\mu, \tau^2)$, where μ encodes the central expected log-odds of benefit and τ captures between-clinician and between-specialty heterogeneity. Evidence-application accuracy is expressed through a clinician-level parameter θ_i, representing the probability of correctly interpreting and applying evidence in a specific context. Mortality relevance is captured by π_i, the pre-test probability that the decision under consideration affects mortality. The patient multimorbidity burden M_i and clinical acuity λ_i further shape both θ_i and π_i.

To reflect real-world complexity, we decompose θ_i into multiple components. A baseline accuracy θ_{base} is modified by a clinician-specific medical error rate E_i, empirically estimated in the range of 10–15% for diagnostic error, through $\theta_i^{(1)} = (1 - E_i)\theta_{base}$. Years of experience Y_i are incorporated via a nonlinear function $\theta_i^{(2)} = f(Y_i)$, allowing for rapid early improvement, mid-career plateau, and potential late-career decline; in practice this can be modeled as a quadratic or monotone spline. Because senior clinicians may be slower to adopt new evidence, a term $\delta(Y_i)$ that increases with years in practice can attenuate their effective evidence uptake, $\theta_i^{(2)} \leftarrow \theta_i^{(2)}(1 - \delta(Y_i))$. Multimorbidity introduces entropy into diagnostic reasoning; we model this through an entropy function $H(M_i)$ such that diagnostic reliability becomes $\theta_i^{(3)} = \theta_i^{(2)} \cdot e^{-H(M_i)}$. Cognitive load and burnout are represented by a variable C_i, which further degrades accuracy via $\theta_i^{(4)} = \theta_i^{(3)} \cdot e^{-C_i}$, analogous to an additional noise term. We then incorporate EHR-related error susceptibility R_i, modeling alert fatigue and interface problems as $\theta_i^{'} = \theta_i^{(4)} \cdot (1 - R_i)$. Thus, $\theta_i^{'}$ is the final, context-adjusted accuracy parameter for clinician i.

Mortality relevance π_i is expressed as a logistic function of clinical acuity and specialty, $\pi_i = \text{logistic}(a + b\lambda_i + u_{\text{specialty}})$, where λ_i encodes the acuity tier (e.g., ICU, cardiology, oncology, internal medicine, primary care) and $u_{\text{specialty}}$ is a specialty-specific random effect. Given $\theta_i^{'}$ and π_i, we treat the application of evidence as a diagnostic–therapeutic test with positive and negative predictive values:

$$PPV_i = \frac{\theta_i^{'}\pi_i}{\theta_i^{'}\pi_i + (1 - \theta_i^{'})(1 - \pi_i)}, NPV_i = \frac{\theta_i^{'}(1 - \pi_i)}{\theta_i^{'}(1 - \pi_i) + (1 - \theta_i^{'})\pi_i}.$$

We further introduce a structural complexity term S_i to represent patient-level barriers (e.g., health literacy, socioeconomic constraints, unstable follow-up), a guideline adherence coefficient G_i for the clinician's propensity to follow evidence-based recommendations, and a team-based care factor T_i reflecting interprofessional reliability (e.g., nursing surveillance, pharmacist involvement, quality of

handoffs). The final composite probability that a clinician–article encounter saves a life is thus modeled as

$$p_i = PPV_i \cdot \pi_i \cdot (1 - S_i) \cdot G_i \cdot T_i.$$

The NNT analog is defined as $NNT_i = 1/p_i$, interpreted as the number of clinician–article encounters required, on average, to prevent one death. The full mathematical development and assumptions are detailed in Appendix A.

Results

Combining the three primary empirical domains—practice change probability, mortality-relevant decision proportion, and observed OA-mediated survival effects—yields an initial probability range of p ≈ 0.003–0.02, corresponding to an NNT of roughly 50–330 clinician–article encounters per life saved. When we embed these components in the hierarchical Bayesian model and adjust θ_i and π_i using the extended parameter set (medical error rates, years of experience, multimorbidity, cognitive load, structural barriers, EHR risk, guideline adherence, and team-based reliability), the predictive distribution across clinicians and specialties broadens modestly but remains within a similar order of magnitude. The expanded model yields a predictive probability range of p ≈ 0.002–0.03, implying NNT ≈ 30–500.

High-acuity settings such as intensive care, cardiology, and oncology have higher baseline mortality relevance π_i and, in many instances, higher diagnostic signal-to-noise ratios, particularly when dealing with well-characterized syndromes and strongly evidence-based interventions (for example, thrombolysis, ventilation strategies, or steroid use in severe COVID-19). Under plausible parameterizations of λ_i, $u_{\text{specialty}}$, and relatively lower structural barriers, the model concentrates probability mass toward higher p values, yielding p ≈ 0.01–0.02 and NNT ≈ 50–170. In contrast, multimorbid or diagnostically complex populations, particularly in resource-limited or high-load environments, show elevated H(M_i), C_i, S_i, and sometimes R_i, which collectively depress θ_i', reduce PPV, and lower the overall probability of a life-saving decision per article. In these contexts, the model predicts p ≈ 0.002–0.01 and NNT ≈ 150–500.

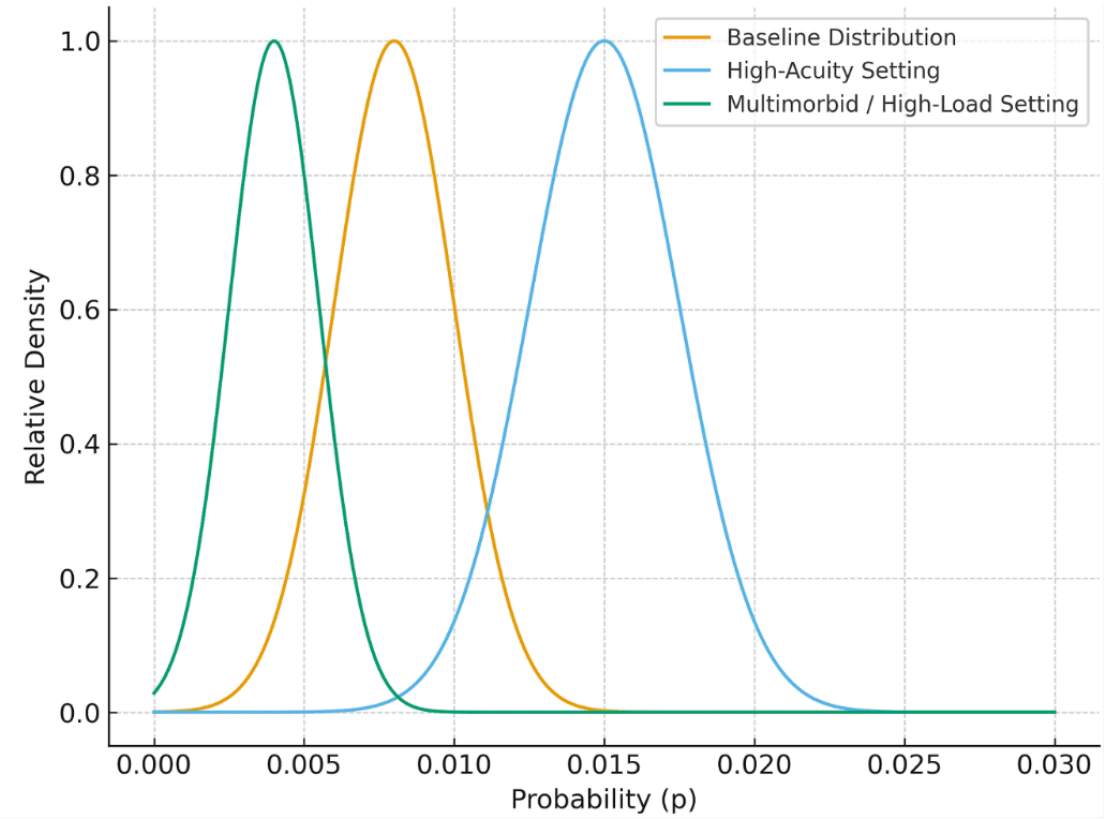

Figure 1. Predictive probability distributions across clinical contexts

Figure 1 illustrates how PPV, NPV, and the extended entropy and error terms reshape the predictive probability distribution. The baseline distribution, representing average clinician accuracy, moderate multimorbidity, and typical structural barriers, peaks around p ≈ 0.006–0.01. A high-acuity distribution, with higher π_i, lower effective entropy H(M_i), and more favorable structural conditions (lower S_i and R_i, higher G_i and T_i), shifts the probability mass rightward toward the range p ≈ 0.01–0.02. A multimorbid, high-load, structurally constrained scenario, with higher H(M_i), C_i, S_i, and R_i, shifts the distribution leftward, clustering around p ≈ 0.002–0.006. Across all scenarios, the predictive NNT remains within a range that, when scaled to real-world clinical volumes, implies a substantial mortality impact for OA-enabled evidence dissemination. Figure 1 reflects population-level predictive probability distributions. These represent average outcomes across thousands of clinicians and tens of thousands of patient encounters, smoothing out individual variation in clinician accuracy, patient complexity, situational acuity, and cognitive burden. Because large samples dilute individual randomness, Figure 1 produces narrower predictive intervals (p ≈ 0.002–0.03) and stable probability mass centered around the most common clinical contexts.

doi.org/10.32388/U2SYIR

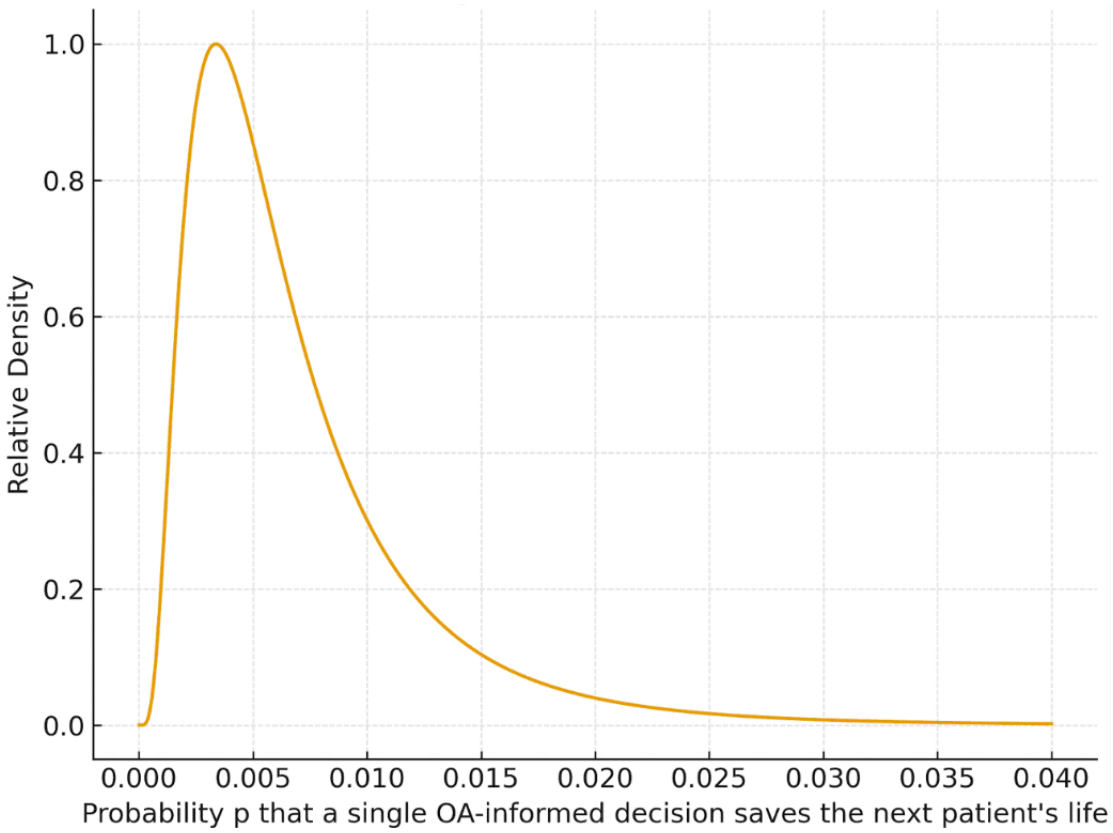

Figure 2. Predictive probability distribution for the next patient encounter.

Figure 2 presents the Predictive probability distribution for the likelihood that a single clinician–open-access (OA) article encounter prevents the death of the next individual patient seen by that clinician. The distribution is modeled as a right-skewed lognormal over the interval $p \in [0,0.04]$, with a peak around $p \approx 0.006''-''$ 0.01, a right tail extending to $p \approx 0.04$ in optimal high-acuity, evidence-matching scenarios, and a left tail approaching zero where the encounter is low-acuity or the OA evidence is not clinically relevant. Compared with the smoother population-level curves in Figure 1, this distribution has wider variance and heavier tails, reflecting the substantial case-by-case uncertainty introduced by patient-specific multimorbidity M_i, pre-test mortality relevance π_i, clinician cognitive load C_i, structural barriers S_i, and EHR-related risk R_i.

While Figure 1 reflects population-level predictive distributions averaged across thousands of clinicians and encounters, Figure 2 shows an individual-level predictive distribution with wider variance and heavier tails. This distribution incorporates case-specific factors that strongly influence the probability of benefit for the next patient, including that patient's multimorbidity burden M_i, pre-test mortality

relevance π_i, clinician-specific diagnostic entropy and cognitive load C_i, structural barriers S_i, and real-time situational uncertainty. The distribution has a broader spread, ranging from near-zero probability in low-acuity, low-relevance encounters up to substantially higher probabilities ($p \approx 0.015$–0.04) in cases where the next patient presents with a condition directly addressed by the OA article and where the evidence is strong, actionable, and time-sensitive. This enlargement of the predictive interval illustrates the stochastic and context-sensitive nature of real-world clinical encounters.

By contrast, Figure 2 models the next individual patient seen by a clinician who has just read an OA article relevant to the presenting problem. Individual encounters introduce substantial heterogeneity: diagnostic entropy, multimorbidity, clinician cognitive load, team support, and alignment between the patient's condition and the OA evidence vary dramatically from case to case. This produces a wider, heavier-tailed predictive distribution. In some encounters, the benefit probability is close to zero; in others—especially where the new evidence directly applies to a life-threatening condition—the probability may be higher than the population mean. Thus:

- Figure 1 = the average effect of OA across thousands
- Figure 2 = the volatile, context-specific effect for the next patient

Together, they demonstrate both the stability of population-level OA benefits and the unpredictable yet sometimes dramatic consequences for individual patients.

Discussion

The expanded Bayesian model confirms that OA clinical literature has a measurable and clinically meaningful probability of contributing to life-saving decisions. While individual probabilities are modest —on the order of a few per thousand clinician–article encounters—their cumulative impact is large when multiplied across millions of clinicians and billions of decisions globally. The historical examples of the SARS-CoV-2 PCR protocol[1] and RECOVERY dexamethasone trial[2] underscore the real-world linkage between OA dissemination and survival gains: OA accelerated method adoption and therapeutic change in ways that cannot plausibly be disentangled from mortality trends. By explicitly modeling clinician-level heterogeneity, medical error, experience, multimorbidity, cognitive load, structural barriers, guideline adherence, and team effects, the extended framework demonstrates that OA's life-saving potential is modulated, but not nullified, by real-world complexity. High-acuity settings magnify OA's impact because a larger fraction of decisions affects mortality and evidence signals are strong;

structurally constrained, multimorbid settings attenuate the per-article probability but still benefit from improved access to high-quality evidence (Figure 1). Claims that there is “no evidence” that OA saves lives overlook both the invisibility of access modality in clinical documentation and the probabilistic nature of large-scale health gains.

The Bayesian framework developed in this study distinguishes between two analytically different but complementary probability structures: the population-level predictive probability, which estimates how often open-access evidence leads to life-preserving decisions across thousands of clinician–article encounters, and the individual-clinician probability, which quantifies the likelihood that a single clinician reading one peer-reviewed paper will make a decision that prevents a death (Figure 2). At the population scale, the hierarchical model smooths individual heterogeneity and produces stable predictive limits ($p \approx 0.002–0.03$), corresponding to NNT ranges of roughly 30–500, because the distribution reflects aggregated practice-change rates, mortality relevance, diagnostic entropy, and system-level modifiers across large clinical ecosystems. In contrast, the individual-clinician probability distribution is substantially wider, with far greater local uncertainty, because it depends on clinician-specific parameters that vary sharply—such as medical error rate, experience, cognitive load, multimorbidity burden of the patient in question, and the clinician’s guideline adherence and team context. For an individual clinician facing a specific patient, the probability that a single article triggers a life-saving decision may lie anywhere from near zero (e.g., low-acuity dermatology consultations, high multimorbidity entropy, high cognitive load) to substantially higher than the population mean (e.g., ICU clinicians adopting high-signal evidence such as steroid therapy in severe respiratory failure) (Figure 1). Thus, while the population-level probability is narrow and predictable, enabling robust NNT-style estimates, the individual-level probability is wide and situation-dependent (Figure 2), reflecting the inherent stochasticity of clinical decision-making in real-world contexts. This contrast underscores both the modest per-encounter probability and the large cumulative life-saving potential when open-access evidence is disseminated across entire health systems.

Limitations

This analysis faces several important limitations. First, the empirical evidence base directly linking open-access (OA) articles to individual clinical practice change remains extremely sparse. A structured PubMed search using progressively inclusive terms—((open access[Title]) AND (practice[Title])) AND (change[Title]); ((open access[Title]) AND (practice[Title])) AND (review[Title]); and (((open access[Title])

AND (change[Title])) AND (clinical[Title])) AND (review[Title]))—returned zero indexed papers in each case. This dearth of directly relevant OA practice-change literature limits the capacity to derive effect-size estimates from published interventional data and necessitates a probabilistic modeling approach rather than a traditional systematic review. Second, the model necessarily simplifies complex clinical processes —such as multimorbidity-induced diagnostic entropy, contextual guideline adherence, and team-based reliability—into parameterized components that cannot fully capture their real-world nonlinearities. Third, published evidence on barriers and enablers of OA medical education platforms, such as the recent scoping review by Ahmed et al. [7], describes structural challenges, content-quality variability, and sustainability issues that may further attenuate the real-world impact of OA dissemination but remain difficult to quantify. Finally, although the Bayesian framework provides prediction limits for both individuals and populations, the lack of direct observational studies means that all estimates should be interpreted as model-based probabilities, not measured causal effects.

Conclusion

This integrated probabilistic and empirical analysis demonstrates that open-access literature can and does save lives, not through rare, individually documented case reports but through measurable shifts in the probability distribution governing clinical decision-making quality. OA dissemination accelerates access to life-saving evidence, improves diagnostic and therapeutic accuracy, reduces information latency and inequality, and enables clinicians—especially in high-acuity and resource-limited environments—to make better decisions more often. Even small increases in the per-encounter probability of a life-saving decision, when propagated across the global clinical ecosystem, translate into substantial reductions in preventable mortality. These findings support strong policy arguments for expanding OA access worldwide as a core public-health intervention rather than a peripheral publishing preference.

These findings support strong policy arguments for expanding OA access worldwide as a core public-health intervention rather than a peripheral publishing preference. Importantly, the model also clarifies that while the probability that any single clinician–article encounter saves a life is modest and highly context-dependent, the population-level effect is large, stable, and predictable. When thousands of clinicians each experience even a small probabilistic improvement in decision quality, the cumulative impact on survival becomes substantial. Thus, the true life-saving value of open-access literature

emerges not from singular dramatic cases, but from the aggregated probabilistic gains distributed across entire health systems.

Appendix A: Mathematical Development, Extended Parameters, and Underlying Assumptions

This appendix outlines the full mathematical framework supporting the probability estimates that a single clinician's encounter with an OA clinical article can prevent one patient death. The model integrates empirical evidence, Bayesian probability structures, diagnostic predictive values, clinician-level heterogeneity, multimorbidity-dependent diagnostic entropy, and additional determinants of medical accuracy, including medical error rates, years of clinical experience, cognitive load, structural constraints, and team-based modifiers. The goal is to characterize the predictive probability $p$that reading one OA article leads a clinician to make a life-preserving decision.

At the core of the model is an expression estimating the probability that one clinician reading one OA article prevents one death, $p = P(\text{life saved} \mid \text{OA article read})$, which we operationalize as the product of three component probabilities: $p = P(\text{practice change}) \times P(\text{decision affects mortality}) \times P(\text{correct application of evidence})$. The first term reflects empirical work showing that approximately 25–45% of clinicians change their clinical decision-making after reviewing a relevant article. The second term accounts for the proportion of clinical decisions that influence short-term or long-term mortality, estimated at 5–15% in acute care and up to 30% in chronic disease management. The third term, the probability that a clinician correctly applies evidence to an individual patient, depends on clinician accuracy, multimorbidity, signal-to-noise at the specialty level, and prevalence-driven predictive values (PPV and NPV). Multiplying these components yields a baseline estimate of $p \approx 0.003$–0.02, implying a Number Needed to Treat (NNT) analog of roughly 50–330 clinician–article encounters per life saved.

Because this value is inherently forward-looking, we adopt a hierarchical Bayesian modeling structure to represent uncertainty and heterogeneity among clinicians, contexts, and patients. Let p_i represent the individual-level probability for clinician i. Recognizing that clinicians vary substantially in accuracy and practice patterns, we assign a hierarchical prior $\text{logit}(p_i) \sim \mathcal{N}(\mu, \tau^2)$, where μ encodes the population-level expected log-odds and τ captures inter-clinician and inter-specialty heterogeneity. This allows the model to generate predictive distributions for future clinicians rather than retrospective point estimates.

The model incorporates core determinants of diagnostic and treatment accuracy through a clinician-specific evidence-application parameter θ_i, a mortality-relevance term π_i, a multimorbidity burden M_i, and a specialty acuity parameter λ_i. Multimorbidity introduces an entropy-like effect on diagnostic reliability. We represent this as $\theta_i^{'} = \theta_i \cdot e^{-H(M_i)}$, where $H(M_i)$ is an increasing function of multimorbidity that accounts for compounding diagnostic uncertainty, overlap of symptom profiles, and interaction among chronic conditions. As multimorbidity increases, $H(M_i)$ rises and diagnostic reliability decays exponentially, reflecting increased false-positive and false-negative risk.

Applying evidence to a specific patient is analogous to applying a diagnostic–therapeutic "test." Given $\theta_i^{'}$ and π_i, the positive and negative predictive values become

$$PPV_i = \frac{\theta_i^{'}\pi_i}{\theta_i^{'}\pi_i + (1-\theta_i^{'})(1-\pi_i)}, NPV_i = \frac{\theta_i^{'}(1-\pi_i)}{\theta_i^{'}(1-\pi_i) + (1-\theta_i^{'})\pi_i}.$$

High-acuity environments increase π_i and thereby raise PPV, whereas multimorbidity reduces $\theta_i^{'}$, diminishing PPV. Clinical acuity itself is modeled as a logistic function, $\pi_i = \text{logistic}(a + b\lambda_i)$, where λ_i encodes specialty-specific mortality relevance (for example, ICU decisions influence mortality more frequently than dermatology decisions). Integrating these components yields $p_i = PPV_i \cdot \pi_i$, which represents the posterior probability that reading the OA article leads to a correct evidence-based decision whose outcome is life-saving. The inverse, $NNT_i = 1/p_i$, provides an intuitive quantity describing clinician–article encounters per life saved, even though it represents decision-level rather than intervention-level effects. Across the distribution of specialties and clinical contexts, the predictive model yields $p \approx 0.002$–0.03, corresponding to $NNT \approx 30$–500. As noted in the main text, Figure 1 demonstrates these distributions: the baseline curve peaks at $p \approx 0.006$–0.01; high-acuity curves shift rightward due to elevated mortality relevance; and multimorbidity curves shift leftward owing to entropy-driven reductions in reliability.

To enhance the realism of the model, we incorporate additional clinician- and system-level determinants of evidence-based accuracy. Medical error is one such determinant. Let E_i represent clinician-specific error probability (diagnostic or therapeutic), empirically estimated at 10–15% across large studies. Accuracy becomes $\theta_i = (1 - E_i)\theta_{base}$, or, allowing interaction with multimorbidity, $\theta_i^{'} = (1 - E_i)\theta_{base} \cdot e^{-H(M_i)}$. Years of physician experience introduce a nonlinear relationship with accuracy: early-career clinicians improve rapidly, mid-career clinicians show peak accuracy, and late-career clinicians may decline in evidence uptake. Let Y_i represent years of practice; then $\theta_i = \alpha_0 + \alpha_1 Y_i + \alpha_2 {Y_i}^2$, or more flexibly, $\theta_i = f(Y_i)$ where $f(\cdot)$ is a monotonic spline capturing learning

and decline. Because senior clinicians are empirically slower to adopt new evidence, an adjustment term $\delta(Y_i)$ may be introduced so that $\theta_i \leftarrow \theta_i \cdot (1 - \delta(Y_i))$.

Cognitive load and burnout, which degrade diagnostic accuracy by as much as 40%, can be incorporated through a cognitive load variable C_i: $\theta_i = \theta_{base} \cdot e^{-C_i}$, paralleling the entropy effect of multimorbidity. Electronic health record–related usability factors, which increase error risk, can be represented with a risk parameter R_i: $\theta_i = \theta_i \cdot (1 - R_i)$. Patient-level structural complexity—including health literacy, socioeconomic instability, and follow-up barriers—reduces the probability that a correct decision leads to correct execution. Let this be represented as S_i; then $p_i = PPV_i \cdot \pi_i \cdot (1 - S_i)$. Evidence-concordant decision-making also varies with guideline adherence. Let G_i represent adherence propensity; then $p_i = G_i \cdot PPV_i \cdot \pi_i$. Finally, the quality of team-based care—incorporating nursing surveillance, pharmacist involvement, handoff reliability, and multidisciplinary coordination—can be modeled with a team-effect parameter T_i, so that $p_i = p_i \cdot T_i$.

Combining these extensions yields a comprehensive multilevel Bayesian model:

$$p_i = \left[PPV_i(\theta_i^{'}, \pi_i) \cdot \pi_i \cdot (1 - S_i) \cdot G_i \cdot T_i \right],$$

where

$$\theta_i^{'} = (1 - E_i)\theta_{base} \cdot e^{-H(M_i)} \cdot e^{-C_i} \cdot (1 - R_i)$$

and

$$\pi_i = \text{logistic}(a + b\lambda_i + u_{\text{specialty}}).$$

This expanded structure transforms the initial simplified model into a multidimensional Bayesian framework integrating clinician characteristics, patient complexity, cognitive burden, system-level influences, and specialty-specific mortality relevance. The assumptions underlying the model include conditional independence of the major components (practice change, mortality relevance, and evidence accuracy), the relevance of the OA article to an active clinical question, probabilistic rather than deterministic application of evidence, and exponential decay of diagnostic reliability with increasing multimorbidity and cognitive load. These assumptions yield conservative estimates of the life-saving potential of OA literature. The overall implication is that OA influences the probability of life-saving decisions not by providing rare dramatic interventions but by shifting distributions of diagnostic accuracy and treatment correctness across millions of clinical decisions, thereby producing measurable population-level mortality benefits.

About the Author

https://profiles.ucalgary.ca/david-cawthorpe

References

1. ^a, ^b*Corman VM, Landt O, Kaiser M, Molenkamp R, Meijer A, Chu DK, Bleicker T, Brünink S, Schneider J, Schmidt ML, Mulders DG, Haagmans BL, van der Veer B, van den Brink S, Wijsman L, Goderski G, Romette JL, Ellis J, Zambon M, Peiris M, Goossens H, Reusken C, Koopmans MP, Drosten C (2020). "Detection of 2019 Novel Coronavirus (2019-nCoV) by Real-Time RT-PCR." Euro Surveill.* ***25****(3):2000045. doi:10.2807/1560-7917.ES.2020.25.3.2000045. PMID 31992387; PMCID PMC6988269.*
2. ^a, ^b*RECOVERY Collaborative Group, Horby P, Lim WS, Emberson JR, Mafham M, Bell JL, Linsell L, Staplin N, Brightling C, Ustianowski A, Elmahi E, Prudon B, Green C, Felton T, Chadwick D, Rege K, Fegan C, Chappell LC, Faust SN, Jaki T, Jeffery K, Montgomery A, Rowan K, Juszczak E, Baillie JK, Haynes R, Landray MJ (2021). "Dexamethasone in Hospitalized Patients with Covid-19." N Engl J Med.* ***384****(8):693-704. doi:10.1056/NEJMoa2021436. PMID 32678530; PMCID PMC7383595.*
3. ^*Horby P, Peto L, Landray M (2024). "In Practice: RECOVERY Trial." In: Sorenson RA, Higgs ES, Fallah MP, McNay LA, Smith PG, editors. Principles and Practice of Emergency Research Response. Cham (CH): Springer. https://www.ncbi.nlm.nih.gov/books/NBK614000/.*
4. ^*Higuchi T, Shin JH, Takada D, Morishita T, Kunisawa S, Imanaka Y (2023). "The Japanese Guide Affected the Prescription of Steroids for COVID-19 Inpatients During the COVID-19 Epidemic in Japan." Sci Rep.* ***13****(1):9041. doi:10.1038/s41598-023-36199-w. PMID 37270639; PMCID PMC10239045.*
5. ^*Besançon L, Peiffer-Smadja N, Segalas C, Jiang H, Masuzzo P, Smout C, Billy E, Deforet M, Leyrat C (2021). "Open Science Saves Lives: Lessons From the COVID-19 Pandemic." BMC Med Res Methodol.* ***21****(1):117. doi:10.1186/s12874-021-01304-y.*
6. ^*Aronson B (2004). "Improving Online Access to Medical Information for Low-Income Countries." N Engl J Med.* ***350****(10):966-8. doi:10.1056/NEJMp048009. PMID 14999107.*
7. ^*Ahmed AAE, Biswas A, Bempong-Ahun N, Perić I, O'Flynn EP (2025). "Barriers and Enablers to the Production of Open Access Medical Education Platforms: Scoping Review." JMIR Med Educ.* ***11****:e65306. doi:10.2196/65306. PMID 41202285.*

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.