Supplementary Material for Evidence Readiness Assessment for Drug Repurposing Using Data Quality, Safety Signals, Network Structure, and Bayesian Uncertainty

Francis Osei[1]

francis@bayezian.com

1 Supplementary Methods

1.1 Model constants and parameter trace

The evidence-readiness model used a fixed set of constants for the main analysis. These constants governed the concentration of the literature-derived Beta prior, the strength of the graph-derived likelihood contribution, the adverse-evidence penalty, and the safety-overlap adjustment. Table S1 reports the parameter values used in the main analysis. The concentration function was

Table S1. Model constants used in the main evidence-readiness analysis.

Parameter	Symbol	Value	Role in the model
Maximum concentration	$c_{\max}$	200	Upper bound for the article-count contribution to the Beta prior concentration.
Concentration rate	τ	25	Rate parameter controlling the increase in prior concentration as usable article count increases.
Likelihood strength	λ	50	Fixed pseudo-count strength assigned to the graph-derived likelihood contribution.
Adverse-evidence weight	w_A	2	Weight applied to adverse literature in the penalised semantic evidence score.
Safety dampening coefficient	κ	0.5	Coefficient applied to the safety-overlap term γ in the final semantic evidence adjustment.

$$c(M) = c_{\max}\left(1 - \exp\left(-\frac{M}{\tau}\right)\right),$$

where M denotes the number of usable classified articles. The main analysis used $c_{\max} = 200$ and $\tau = 25$. Graph evidence was represented as a Beta likelihood contribution with strength $\lambda = 50$.

The adverse-evidence coefficient $w_A = 2$ was chosen to place greater penalty on adverse semantic evidence than on a matched count of therapeutic evidence. This conservative choice reduces false-positive readiness signals when the retrieved

literature contains harm, worsening, or clinically relevant adverse association signals. The safety coefficient $\kappa = 0.5$ was chosen as a bounded dampening term. It reduces readiness in the presence of safety overlap while avoiding a binary exclusion rule, since safety overlap does not establish causation.

1.2 Sensitivity analysis design

A one-at-a-time perturbation analysis was conducted for the five main constants. Each parameter was varied at $0.5\times$ and $1.5\times$ its baseline value, with all other parameters held at baseline. Trial records, terminology mappings, PubMed/PMC article classifications, safety-overlap scores, and graph features were fixed during perturbation.
A scenario was classified as stable if it met all three criteria shown in Table S2.

Table S2. Pre-specified stability criteria for parameter perturbation scenarios.

Diagnostic	Threshold	Interpretation
Spearman rank correlation	≥ 0.95	Rank order remains close to the baseline ordering.
Median absolute posterior-mean shift	≤ 0.05	Posterior means show limited central displacement.
Unchanged quality flags	$\geq 90\%$	Rule-based evidence categories remain stable for most pairs.

2 Supplementary Results

2.1 Evidence coverage

Coverage tiers were used to distinguish drug-disease pairs with complete Bayesian audit coverage from pairs with partial evidence support. Table S3 reports the number of pairs, unique drugs, and unique diseases in each tier.

Table S3. Evidence coverage tiers across audited drug-disease pairs.

Coverage tier	Drug-disease pairs	Unique drugs	Unique diseases
Literature tier	15	9	13
Full Bayesian	13	9	10
Bayesian without safety	12	5	12
Graph tier	1	1	1

2.2 Quality-flag definitions

Quality flags identify the main interpretation constraint for each pair. These flags were assigned after integration of terminology evidence, literature composition, safety overlap, graph support, posterior uncertainty, and evidence-readiness score.

Table S4. Quality-flag definitions used in the evidence-readiness assessment.

Quality flag	Definition
High evidence quality	Strong multi-domain support, limited adverse burden, limited retrieval noise, acceptable posterior uncertainty.
Moderate evidence quality	Evidence support present, but residual uncertainty, sparsity, or mixed signal requires expert review.
Insufficient evidence	Too little usable evidence for robust interpretation.
Safety-conflicted evidence	Safety-overlap signal and adverse-evidence burden create material conflict.
Safety-concerning	Safety-overlap signal present; review required before prioritisation.
Literature-conflicted	Therapeutic and adverse literature signals both contribute to interpretation.
Literature noise dominated	Irrelevant retrieved articles dominate the usable literature set.
Terminology uncertainty	Drug or disease mapping uncertainty limits evidence interpretation.

2.3 Parameter sensitivity

Table S5 reports the perturbation results for all model constants. Eight of ten scenarios satisfied the pre-specified stability rule. Perturbations of $c_{\max}$, τ, λ, and w_A retained high rank agreement with the baseline analysis. Perturbation of the safety coefficient κ produced greater rank movement, reflecting the direct role of safety overlap in the final semantic evidence adjustment.

Table S5. Sensitivity analysis for model constants.

Perturbed parameter	Multiplier	Value	Pairs	Spearman	Top-10 retained	Median rank shift	Max rank shift	Median $\lvert\Delta\mu\rvert$	Stable
w_A	0.5	1.00	41	0.9613	9	1.0	12	0.0212	Yes
w_A	1.5	3.00	41	0.9890	9	1.0	6	0.0121	Yes
$c_{\max}$	0.5	100.00	41	0.9969	10	0.0	2	0.0231	Yes
$c_{\max}$	1.5	300.00	41	0.9995	10	0.0	1	0.0114	Yes
λ	0.5	25.00	41	0.9970	10	0.0	3	0.0171	Yes
λ	1.5	75.00	41	0.9988	10	0.0	2	0.0124	Yes
κ	0.5	0.25	41	0.9329	7	2.0	13	0.0528	No
κ	1.5	0.75	41	0.9249	8	2.0	15	0.0528	No
τ	0.5	12.50	41	0.9969	10	0.0	3	0.0059	Yes
τ	1.5	37.50	41	0.9983	10	0.0	2	0.0069	Yes

$\lvert\Delta\mu\rvert$ denotes the absolute posterior-mean shift relative to the baseline run. Top-10 retained denotes the number of baseline top-10 ranked pairs also present in the perturbed top-10 set.

2.4 Pair-level evidence-readiness audit

Table S6 reports a compact pair-level audit table. The table includes coverage tier, usable literature count, adverse and irrelevant proportions, safety-overlap score, posterior mean, credible-interval width, evidence-readiness score, and quality flag. These fields were selected to show the evidence basis and the main interpretation constraint without reproducing the full computational ledger.

Table S6: Compact pair-level evidence-readiness audit.

Drug–disease pair	Coverage	Articles	Adverse	Noise	γ	μ_{post}	CI width	ERS	Quality flag
aspirin–colorectal cancer	Literature tier	39	0.026	0.179	0.750	0.473	0.135	61.3	Terminology uncertainty
aspirin–pre-eclampsia	Full Bayesian	40	0.000	0.225	0.200	0.708	0.122	87.9	High evidence quality
azithromycin–albuminuria	Bayesian without safety	5	0.000	0.800	0.200	0.366	0.199	67.1	Literature noise dominated
azithromycin–bronchiectasis	Full Bayesian	40	0.000	0.250	0.850	0.449	0.134	76.6	Safety-concerning
azithromycin–cystic fibrosis	Bayesian without safety	40	0.075	0.450	0.700	0.278	0.120	72.5	Moderate evidence quality
azithromycin–frailty	Bayesian without safety	3	0.000	0.667	0.200	0.436	0.223	63.8	Moderate evidence quality
azithromycin–venous thrombosis	Bayesian without safety	16	0.250	0.563	0.800	0.171	0.121	65.2	Safety-conflicted evidence
colchicine–atherosclerosis	Bayesian without safety	40	0.025	0.200	0.200	0.665	0.127	86.0	High evidence quality
colchicine–coronary artery disease	Full Bayesian	40	0.050	0.250	0.750	0.406	0.132	78.4	Safety-concerning
colchicine–heart failure	Bayesian without safety	40	0.000	0.500	0.800	0.344	0.128	73.1	Safety-concerning
colchicine–myocardial infarction	Bayesian without safety	40	0.050	0.200	0.750	0.425	0.133	77.6	Safety-concerning
dexamethasone–acute myeloid leukaemia	Literature tier	40	0.050	0.500	0.750	0.288	0.122	56.4	Terminology uncertainty
dexamethasone–covid-19	Full Bayesian	40	0.075	0.100	0.850	0.416	0.132	79.2	Safety-concerning
dexamethasone–multiple myeloma	Full Bayesian	40	0.000	0.225	0.850	0.459	0.134	78.1	Safety-concerning
hydroxychloroquine–covid-19	Full Bayesian	40	0.225	0.275	0.750	0.147	0.095	77.0	Safety-conflicted evidence
hydroxychloroquine–lupus erythematosus	Literature tier	40	0.025	0.175	0.850	0.448	0.134	59.9	Terminology uncertainty
hydroxychloroquine–melanoma	Full Bayesian	40	0.100	0.300	0.200	0.425	0.133	85.7	High evidence quality
hydroxychloroquine–pancreatic cancer	Literature tier	40	0.050	0.250	0.700	0.417	0.132	60.8	Terminology uncertainty
hydroxychloroquine–prostate cancer	Literature tier	8	0.250	0.250	0.200	0.244	0.162	60.4	Terminology uncertainty
hydroxychloroquine–rheumatoid arthritis	Literature tier	36	0.083	0.167	0.850	0.377	0.132	59.4	Terminology uncertainty
ivermectin–covid-19	Full Bayesian	40	0.150	0.575	0.700	0.124	0.088	73.8	Moderate evidence quality
ivermectin–malaria	Bayesian without safety	40	0.025	0.225	0.800	0.435	0.133	74.5	Safety-concerning
ivermectin–strongyloidiasis	Literature tier	38	0.026	0.132	0.800	0.480	0.135	61.2	Terminology uncertainty
levodopa–cardiovascular diseases	Bayesian without safety	40	0.300	0.500	0.200	0.118	0.086	79.4	Literature-conflicted
levodopa–depression	Literature tier	31	0.194	0.484	0.850	0.134	0.095	54.1	Terminology uncertainty

Continued on next page

Drug–disease pair	Coverage	Articles	Adverse	Noise	γ	μ_{post}	CI width	ERS	Quality flag
levodopa–sarcopenia	Bayesian without safety	13	0.000	0.615	0.800	0.329	0.159	61.5	Safety-concerning
levodopa–tinnitus	Graph tier	0	NA	0.000	NA	0.481	0.268	62.9	Insufficient evidence
metformin–chronic kidney disease	Literature tier	40	0.050	0.275	0.800	0.383	0.131	58.9	Terminology uncertainty
metformin–colorectal cancer	Literature tier	40	0.025	0.250	0.700	0.454	0.134	61.0	Terminology uncertainty
metformin–endometrial cancer	Literature tier	40	0.025	0.200	0.200	0.670	0.126	69.4	Terminology uncertainty
metformin–mania	Bayesian without safety	35	0.029	0.371	0.200	0.527	0.137	82.8	High evidence quality
metformin–mitochondrial diseases	Full Bayesian	37	0.378	0.162	0.200	0.126	0.090	85.4	Literature-conflicted
metformin–non-alcoholic fatty liver disease	Full Bayesian	40	0.025	0.425	0.850	0.340	0.127	74.8	Safety-concerning
metformin–uterine cervical dysplasia	Bayesian without safety	8	0.125	0.500	0.200	0.298	0.172	74.7	Moderate evidence quality
rituximab–rheumatoid arthritis	Literature tier	40	0.025	0.125	0.750	0.500	0.134	62.1	Terminology uncertainty
sirolimus–kidney transplant rejection	Literature tier	38	0.105	0.342	0.850	0.272	0.120	56.8	Terminology uncertainty
sirolimus–lymphangioleiomyomatosis	Full Bayesian	40	0.050	0.125	0.850	0.437	0.133	78.9	Safety-concerning
sirolimus–tuberous sclerosis	Full Bayesian	40	0.000	0.200	0.800	0.485	0.134	79.0	Safety-concerning
statins–Alzheimer's disease	Literature tier	40	0.075	0.375	0.700	0.319	0.125	58.7	Terminology uncertainty
statins–sepsis	Literature tier	40	0.050	0.375	0.200	0.481	0.134	66.4	Terminology uncertainty
thalidomide–multiple myeloma	Full Bayesian	40	0.075	0.050	0.850	0.437	0.133	80.0	Safety-concerning

3 Supplementary Interpretation

The sensitivity analysis indicates that the posterior ranking and evidence-readiness categories were stable under perturbation of the prior concentration parameters, the graph likelihood strength, and the adverse-evidence weight. The safety coefficient had greater rank influence because the safety-overlap term enters the final semantic evidence score before Bayesian updating. This supports the main manuscript interpretation that safety overlap should remain a separate audit dimension rather than be absorbed into a single posterior score.