



[Evaluating a Rule-Based Gut Resistome-to-Antibiotic Guidance Framework: Performance, Failure Modes and a Revised Decision Rule]

Seshu Kumari K^{*1}, Dr. Munesh Kumar Sharma^{*2}

^{*1}Phd Scholar, Department of Microbiology (Under Faculty of Medicine), Amaltas University, Dewas, Madhya Pradesh, India

^{*2}Professor and Guide, Department of Microbiology (Under Faculty of Medicine), Amaltas University, Dewas, Madhya Pradesh, India

ABSTRACT

A companion paper specified the architecture of a six-stage computational framework that converts gut-bacterial resistome content into class-level antibiotic guidance. Here we report the quantitative behaviour of that framework when executed over a reference panel of ten gut-relevant bacterial strains and six curated resistance determinants. Seven determinant-present calls were made across the panel (mean burden 0.7 determinants per strain), generating 96 guidance decisions across sixteen clinical antimicrobial classes. Under an original binary decision rule, seven decisions correctly recommended avoidance and eighty-nine asserted that an agent was a preferred candidate; systematic checking against established clinical microbiology showed several of the latter to be false, including a recommendation of cephalosporins for an intrinsically resistant enterococcus and aminoglycosides for obligate anaerobes. A revised three-state rule, supported by an independently compiled intrinsic resistance register, removed every identified error and reduced the number of positive-suitability claims from eighty-nine to zero. Standard descriptive measures of resistome composition — Shannon diversity, ordination, co-occurrence networks, association testing and classifier accuracy — were computed and shown to be uninformative or statistically inapplicable at this panel size, a finding we report as a methodological result rather than a limitation to be glossed over. A forty-community synthetic simulation confirmed that the underlying burden and diversity measures respond in the expected direction to specified compositional differences, while also exposing an internal inconsistency in the permutational test used to compare community dissimilarity. We conclude that resistome profiling of this kind can support avoidance decisions but cannot, on present evidence, support any positive claim of suitability, and we set out the conditions that would need to be satisfied before that could change.

KEYWORDS: antimicrobial resistance; gut resistome; antibiotic selection guidance; clinical decision support; intrinsic resistance; synthetic community simulation; statistical validity; antibiotic stewardship.

Received
18/07/2026

Accepted
20/08/2026

Published
26/09/2026

Article Type
Research Paper



1. Introduction

The gut microbiota is now well established as a reservoir of antimicrobial resistance determinants independent of an individual's own antibiotic history [1], and the global burden of bacterial antimicrobial resistance is estimated at well over a million attributable deaths annually [2], with a substantial additional burden documented at regional scale [3]. Resistome profiling — the systematic detection of resistance determinants in a genome or metagenome against curated reference databases such as CARD [4] — is technically mature. What remains largely unaddressed is the step that would make such profiling clinically actionable: converting a determinant panel into a decision about which antimicrobial class to avoid, and, just as importantly, deciding what may legitimately be inferred when no determinant is found.

A companion design paper specified a six-stage computational framework intended to perform exactly this conversion, built around a small, exhaustively inspectable reference panel of ten gut-relevant bacterial strains and six well-characterised resistance determinants, an ontology-to-clinical-class mapping, and a guidance rule governing how the presence and absence of a determinant are to be interpreted. This paper reports what happened when that framework was executed. We report the resistome content of the panel, the full set of descriptive analytics the framework's sixth stage computes, the guidance decisions generated by an original binary rule and by a revised three-state rule, the clinically erroneous recommendations the revision removed, an internal consistency check against the literature from which the panel was curated, and a synthetic community simulation used to test whether the framework's summary measures respond in the expected direction to a specified compositional difference. Where a result is an artefact of panel size or of the synthetic construction rather than a finding about gut biology, we say so explicitly, because distinguishing the two is itself among the principal contributions of this evaluation.

2. Related Work

Antimicrobial stewardship programmes seek to optimise antimicrobial use to improve patient outcomes while limiting the selection of resistance [5], and standardised definitions of multidrug, extensive drug and pandrug resistance [6] give the vocabulary in which the clinical consequences of a wrong recommendation must be expressed. Country-level resistome surveys show that gut determinant content tracks antibiotic consumption patterns closely [7], and structural studies of resistance-gene distribution across host phylogeny show that determinant carriage is non-random [8], both of which motivate treating a curated, taxonomically representative panel as an adequate testbed for decision logic before that logic is applied to real specimens. The specific hazard evaluated in this paper — that absence of a determinant is systematically misread as evidence of susceptibility — has, to our knowledge, not previously been quantified for a resistome-to-guidance pipeline, although the underlying microbiological fact that intrinsic resistance is common and cannot be detected by gene counting alone is well established for individual organism groups, including the anaerobic Gram-negative bacilli and the enterococci considered here [9].

3. Materials and Methods

Full architectural detail is given in the companion design paper; this section summarises only what is necessary to interpret the results below. The framework was executed over a panel of ten reference gut-bacterial strains and six curated resistance determinants (Table I), implemented in Python using pandas, NumPy, SciPy, scikit-learn, NetworkX, Matplotlib and seaborn. Determinant calls were curated from the primary literature describing each strain rather than generated by sequence annotation; this is stated here because it constrains the interpretation of every result in Section 4, particularly Section 4.7.



Table I. Resistome burden and determinant content by strain.


Strain	Burden	Determinant(s) present	Ontology class(es)
Enterococcus faecalis V583	2	aac(6')-II, vanB operon	Aminoglycoside, glycopeptide
Bifidobacterium longum NCC2705	1	tetW	Tetracycline
Escherichia coli Nissle 1917	1	AmpC β -lactamase	Cephalosporin, penam
Klebsiella pneumoniae MGH78578	1	SHV β -lactamase	Cephalosporin, penam
Bacteroides thetaiotaomicron VPI5482	1	cepA cephalosporinase	Cephalosporin
Bacteroides fragilis NCTC9343	1	cepA cephalosporinase	Cephalosporin
Lactobacillus rhamnosus GG	0	None	—
Faecalibacterium prausnitzii A2165	0	None	—
Clostridioides difficile 630	0	None	—
Akkermansia muciniphila BAA835	0	None	—

Three per-strain descriptive measures — resistome burden, richness and coverage-weighted Shannon diversity — were computed from the strain-by-determinant matrix, together with a principal component ordination, average-linkage hierarchical clustering, a phi-coefficient co-occurrence network, Fisher's exact association testing between determinant carriage and Gram reaction with Benjamini–Hochberg correction [10], and an exploratory random forest classifier of Gram reaction with permutation feature importance. Guidance decisions were generated under two rule variants: an original binary rule (determinant present → avoid; determinant absent → preferred candidate) and a revised three-state rule (determinant present → avoid; determinant absent with recorded intrinsic resistance → avoid; determinant absent with no intrinsic-resistance record → no evidence). Forty synthetic bacterial communities, twenty labelled "healthy" and twenty "dysbiotic", were constructed by the investigator to differ in resistome burden, richness and diversity by specification, and compared by the Mann–Whitney U test and by permutational multivariate analysis of variance on Bray–Curtis dissimilarity.

4. Results

4.1 Resistome Burden and Composition of the Panel

Seven determinant-present calls were made across the ten strains, giving a mean resistome burden of 0.7 determinants per strain (Table I, Fig. 1). Enterococcus faecalis V583 carried two determinants, the highest in the panel; five strains carried one determinant each; four strains — Lactobacillus rhamnosus GG, Faecalibacterium prausnitzii A2165, Clostridioides difficile 630 and Akkermansia muciniphila BAA835 — carried none. Aggregated across the panel, the seven determinant calls produced nine determinant-to-class assignments, since both panel β -lactamases map to two ontology classes each; cephalosporin was the most represented class (four assignments), followed by penam (two), and tetracycline, glycopeptide and aminoglycoside (one each). The pattern reproduces the taxonomy exactly, since the input was curated from sources describing precisely these strain–determinant associations, and should not be read as an independent finding about gut microbial ecology.



4.2 Distribution Across Ontology and Clinical Drug Classes

The strain-by-ontology-class matrix was sparse: of fifty possible cells, nine were occupied. Mapped onward to the sixteen clinical antimicrobial classes, four strains registered a determinant against cephalosporins, one against tetracyclines, one against glycopeptides and one against aminoglycosides; the remaining twelve clinical classes registered no determinant for any strain. This is the clearest single illustration of the framework's central limitation: twelve of sixteen clinical columns were empty throughout, not because the panel organisms are susceptible to those twelve classes, but because the six-determinant panel contains nothing that maps to them.

4.3 Descriptive Analytics: Diversity, Ordination, Co-occurrence and Classification

Coverage-weighted Shannon diversity was non-zero for exactly one strain (*Enterococcus faecalis* V583, at approximately 0.69, equal to $\ln 2$) and zero for every other strain, a result requiring no biological interpretation: a Shannon index computed over determinants takes only the value 0 for a strain with a single determinant and $\ln 2$ for a strain with two determinants of equal weight, so the measure simply relabels which strain carries two determinants rather than one. Principal component analysis of the strain-by-determinant matrix placed 34.8% of variance on the first component and 19.9% on the second (54.7% combined), with *Enterococcus faecalis* separating on the first component and the two *Bacteroides* strains on the second; the four determinant-free strains, whose rows in the matrix are identical, necessarily collapsed onto coincident ordination positions. A phi-coefficient co-occurrence network at a threshold of 0.3 retained a single edge, between the *vanB* operon and the *aac(6')*-II acetyltransferase, both of which occur in exactly one strain; with ten observations and each determinant present in at most two strains, this edge reflects joint presence in a single strain rather than any biologically meaningful linkage. Fisher's exact tests of association between each determinant and Gram reaction returned no result approaching $p = 0.05$ for any determinant (strongest: *cepA* cephalosporinase, $p \approx 0.45$), and no association survived Benjamini–Hochberg correction; at this sample size



the smallest attainable two-sided Fisher p-value exceeds 0.05 regardless of the underlying biology, so non-significance here is a property of the test's applicability, not a negative biological finding. A random forest classifier of Gram reaction from determinant profile returned a training accuracy of 0.90, with the highest permutation importance attributed to the cepA cephalosporinase (≈ 0.27); because the classifier was evaluated without a held-out test set or cross-validation, this figure is a training accuracy only and carries no claim of generalisation.

4.4 Guidance Output Under the Original Binary Rule

The guidance stage generated 96 decisions — six determinant-carrying strains and four determinant-free strains, all ten strains evaluated across sixteen clinical classes. Under the original binary rule, seven decisions returned avoid or use with caution and 89 returned preferred candidate (Fig. 3). Each of the seven avoidance decisions was correct and corresponded to an established resistance of the organism concerned.

4.5 Failure of the Guidance Rule in the Absence of a Determinant

The 89 preferred-candidate decisions were examined individually against standard clinical microbiological references. The framework recommended cephalosporins as a preferred candidate for *Enterococcus faecalis* V583, despite enterococci being intrinsically resistant to all cephalosporins through low-affinity penicillin-binding proteins, one of the most firmly established facts in clinical microbiology. It recommended aminoglycosides as a preferred candidate for both *Bacteroides* strains, despite obligate anaerobes being unable to accumulate aminoglycosides through the oxygen-dependent uptake process the drug class requires. It further recommended fluoroquinolones, macrolides and other classes for organisms in which susceptibility is variable or absent (Table II). Every one of these errors shares a single cause: the rule treated a determinant count of zero as evidence of susceptibility, when it is only evidence that the six-determinant panel contains nothing mapping to that class. Two further errors arose within classes the panel does represent: both *Enterobacterales* strains were recommended penicillins despite carrying a β -lactamase mapping to the penam ontology class, traced to an implementation defect in which the guidance stage counted determinants per clinical class after the ontology mapping had already collapsed, losing the penam assignment before the count was taken — a defect distinct in kind from the conceptual error above, and one identified only because the panel was small enough to check by hand.

Table II. Selected clinically erroneous guidance produced by the original binary rule.

Strain	Class recommended	Rule output	Established position
<i>Enterococcus faecalis</i> V583	Cephalosporins	Preferred candidate	Intrinsically resistant (low-affinity PBPs)
<i>Bacteroides fragilis</i> NCTC9343	Aminoglycosides	Preferred candidate	Intrinsically resistant (no oxidative uptake)
<i>Bacteroides thetaiotaomicron</i> VPI5482	Aminoglycosides	Preferred candidate	Intrinsically resistant (no oxidative uptake)
<i>Bacteroides fragilis</i> NCTC9343	Fluoroquinolones	Preferred candidate	Variable, frequently non-susceptible
<i>Bacteroides thetaiotaomicron</i> VPI5482	Macrolides	Preferred candidate	Frequently non-susceptible
<i>Escherichia coli</i> Nissle 1917	Penicillins	Preferred candidate	AmpC confers aminopenicillin resistance





Klebsiella pneumoniae MGH78578	Penicillins	Preferred candidate	Chromosomal SHV confers ampicillin resistance
Enterococcus faecalis V583	Carbapenems	Preferred candidate	Reduced activity, not first line

4.6 Guidance Output Under the Revised Three-State Rule

Applying the revised three-state rule together with the intrinsic resistance register removed every error listed in Table II. Of the 96 decisions, 7 (7.3%) reflected direct determinant evidence and recommended avoidance; 23 (24.0%) had no panel determinant but a recorded intrinsic resistance and were likewise resolved to avoidance; and 66 (68.8%) had neither a panel determinant nor a recorded intrinsic resistance and were reported as absence of evidence, generating no positive claim (Fig. 2). The number of decisions asserting that an agent is a preferred candidate fell from 89 under the original rule to zero under the revised rule (Fig. 3).

Figure 2. Distribution of the 96 guidance decisions under the revised three-state rule





4.7 Internal Consistency Check

Framework output was compared against the phenotype described in the source publication for each of the nine strains with a documented resistance phenotype (Table III); agreement was complete (9 of 9, 100%). Because the determinant calls analysed throughout this paper were curated by the investigator from these same source publications rather than generated by independent sequence annotation (Section 3), this figure tests transmission of input to output through the mapping and guidance stages — that is, whether the framework preserves and correctly acts on the information it was given — and carries no evidential weight regarding predictive accuracy on unseen genomes. We report it as a consistency check, not a validation, and flag any contrary framing in secondary summaries of this work as an error.

Table III. Internal consistency check against source literature.


Strain	Expected phenotype	Determinant(s) emitted	Consistent
Enterococcus faecalis V583	Vancomycin resistance (VanB); intrinsic aminoglycoside resistance	vanB operon; aac(6')-II	Yes
Bifidobacterium longum NCC2705	Tetracycline resistance (TetW)	tetW	Yes
Klebsiella pneumoniae MGH78578	SHV β -lactamase	SHV β -lactamase	Yes
Escherichia coli Nissle 1917	AmpC cephalosporin resistance	AmpC β -lactamase	Yes
Bacteroides fragilis NCTC9343	cepA cephalosporinase	cepA	Yes
Lactobacillus rhamnosus GG	No acquired determinants	None	Yes
Akkermansia muciniphila BAA835	No acquired determinants	None	Yes
Clostridioides difficile 630	Low acquired determinant load	None	Yes
Concordance			9 of 9 (100%)

4.8 Synthetic Community Simulation



Forty synthetic communities were constructed, twenty designated "healthy" and twenty "dysbiotic" (Fig. 4). Coverage-weighted resistome burden was 0.572 in the healthy group and 0.801 in the dysbiotic group; resistome richness was exactly 3 in every healthy community and exactly 6 in every dysbiotic community, with no within-group variation in either arm; Shannon diversity was approximately 1.03 and 1.71 respectively. The Mann–Whitney U test returned $p < 0.0001$ for all three measures. None of these figures is an empirical finding about real gut communities: the two groups were assembled by the investigator to differ in exactly these respects, and the complete absence of within-group variation in richness — a pattern no set of real biological specimens would be expected to produce — is itself the signature of synthetic rather than sampled data. Permutational multivariate analysis of variance on Bray–Curtis dissimilarity across all forty communities returned a pseudo-F of 0.452 with $p = 0.001$. A pseudo-F below unity indicates that between-group dissimilarity was smaller than within-group dissimilarity, which is internally inconsistent with a significant permutational result and indicates a defect either in the permutation scheme or in the construction of the dissimilarity matrix; we report this as a finding about the implementation rather than suppressing it. The simulation establishes only that the burden and diversity measures of Section 3 respond in the expected direction when community composition is altered in a specified way; it establishes nothing about real gut communities.

Figure 4. Synthetic community simulation ($n = 40$ communities; Mann-Whitney $p < 0.0001$ for all three measures). All differences are a restatement of the investigator-specified construction, not an empirical finding.



4.9 Summary of Principal Findings

(i) The framework executed end to end over ten strains and six determinants, producing 96 guidance decisions, with correct transmission of input to output in all nine testable cases. (ii) All seven avoidance decisions corresponded to established resistances. (iii) The original binary rule failed systematically in the absence of a determinant, producing recommendations contradicted by established microbiological fact; the revised three-state rule, supported by the intrinsic resistance register, removed every identified error at the cost of asserting suitability for no class under any circumstance. (iv) Standard descriptive measures of resistome composition — diversity, ordination, co-occurrence and association testing — are, at this panel size, largely uninformative or beyond the limit of their statistical applicability, and classifier accuracy without a held-out set is a training accuracy only. (v) The 100% internal consistency figure is an artefact of the circular relationship between curated input and curated source, not a predictive validation. (vi) The synthetic simulation reproduces its own specification and additionally exposes an internal inconsistency in the permutational test.

5. Discussion

5.1 Absence of a Determinant Is Not Evidence of Susceptibility



The single most consequential finding of this evaluation is not that the framework worked, but the specific way in which its first implementation failed. A determinant count of zero within a six-gene panel means only that the panel contains nothing mapping to that clinical class; converting that into a positive recommendation produced clinically false guidance in every affected case examined, including two instances — cephalosporins for an enterococcus, aminoglycosides for obligate anaerobes — that contradict facts no practising clinical microbiologist would dispute. This is not a defect specific to the small panel used here; it is a structural property of any determinant-counting procedure applied to a necessarily incomplete reference panel, and it will recur at any panel size unless absence is explicitly represented as absence of evidence rather than evidence of absence.

5.2 Statistical Validity at Small Panel Sizes

Several of the descriptive analyses reported in Section 4.3 were, by construction, incapable of returning an informative result at this panel size: the smallest attainable Fisher exact p-value with ten strains divided across two groups exceeds the conventional 0.05 threshold regardless of the underlying biology, and a Shannon index computed over a maximum of two determinants per strain can take only two distinct values. We report these not as failures of the framework but as a general caution for the resistome literature: applying community-ecology statistics developed for large, continuously varying assemblages to small, sparse determinant panels can manufacture an appearance of either signal or its absence that reflects the arithmetic of the test rather than the biology under study, and both directions of error are misleading in a clinical context.

5.3 The Concordance Figure Is Not a Validation

The 100% internal consistency figure of Section 4.7 will, if read carelessly, look like strong evidence that the framework predicts resistance phenotype accurately. It is not evidence of that at all: the determinant calls used as input were curated from the same publications against which the output was checked, so the check demonstrates only that information is transmitted through the mapping and guidance stages without corruption. A genuine validation requires determinant calls generated independently of the comparison phenotype, most straightforwardly by executing sequence annotation against genomes with linked, independently determined antimicrobial susceptibility results.

5.4 Limitations

The limitations of the underlying framework are set out in the companion design paper and are not repeated in full here; three bear directly on the results above. First, because determinant calls were curated rather than computed, none of the results in this paper demonstrates anything about the framework's behaviour on tool-generated sequence annotation output, which remains untested. Second, the synthetic community simulation of Section 4.8 was constructed by specification and is reported as a check on the direction of the burden and diversity measures, not as an empirical finding about gut community structure, and the pseudo-F anomaly identified there indicates an implementation defect requiring correction before the permutational test is used again. Third, the statistical analyses throughout this paper are conducted at or beyond the limit of their applicability given the panel size, a constraint inherent to the exhaustively-inspectable design defended in the companion paper and traded off deliberately against statistical power.

5.5 Implications for Resistome-Informed Stewardship

Taken together, these results support a conservative conclusion. Resistome data of the kind modelled here can support a decision to avoid a class when a determinant is detected, and can support avoidance when a documented intrinsic resistance applies, but cannot, on the evidence available from this or any framework built on a finite determinant panel, support a positive claim that an agent is suitable. Four conditions would need to be satisfied before resistome-derived guidance of this kind could responsibly inform prescribing: determinant calls generated by validated sequence



annotation rather than curation; a determinant panel extended to the classes of greatest current clinical concern; an intrinsic resistance register extended well beyond the ten species considered here; and external validation against genomes linked to independently determined susceptibility phenotypes. None of the four is satisfied by the present work. Until they are, the defensible role for resistome profiling of this kind is as a modifier of the prior probability of resistance in empirical prescribing and as an instrument for antimicrobial stewardship surveillance, not as a source of positive prescribing recommendations.

6. Conclusion

Executed over a ten-strain, six-determinant reference panel, the resistome-to-guidance framework described in the companion design paper produced correct avoidance decisions in every case supported by determinant evidence, but produced clinically false positive-suitability recommendations in the absence of a determinant under its original binary rule. A revised three-state rule, informed by an independently compiled intrinsic resistance register, eliminated every identified error at the principled cost of never asserting that an agent is suitable. Standard descriptive resistome statistics were shown to be largely uninformative at this panel size, and an apparently perfect internal consistency figure was shown to reflect circular construction rather than predictive validity. The responsible reading of these results is conservative: resistome profiling can tell a clinician what to avoid when the evidence supports it, and can appropriately register when no positive claim is warranted, but it cannot yet tell a clinician what is safe to use.



References

- [1] W. E. Anthony, C. A. D. Burnham, G. Dantas, and J. H. Kwon, "The gut microbiome as a reservoir for antimicrobial resistance," *The Journal of Infectious Diseases*, vol. 223, no. Suppl. 3, pp. S209–S213, 2021. <https://doi.org/10.1093/infdis/jiaa528>
- [2] C. J. L. Murray et al., "Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis," *The Lancet*, vol. 399, no. 10325, pp. 629–655, 2022. [https://doi.org/10.1016/S0140-6736\(21\)02724-0](https://doi.org/10.1016/S0140-6736(21)02724-0)
- [3] A. Cassini et al., "Attributable deaths and disability-adjusted life-years caused by infections with antibiotic-resistant bacteria in the EU and the European Economic Area in 2015," *The Lancet Infectious Diseases*, vol. 19, no. 1, pp. 56–66, 2019. [https://doi.org/10.1016/S1473-3099\(18\)30605-4](https://doi.org/10.1016/S1473-3099(18)30605-4)
- [4] B. P. Alcock et al., "CARD 2020: antibiotic resistance surveillance with the Comprehensive Antibiotic Resistance Database," *Nucleic Acids Research*, vol. 48, no. D1, pp. D517–D525, 2020. <https://doi.org/10.1093/nar/gkz935>
- [5] P. Dadgostar, "Antimicrobial resistance: implications and costs," *Infection and Drug Resistance*, vol. 13, pp. 3903–3910, 2020. <https://doi.org/10.2147/IDR.S234610>
- [6] A. P. Magiorakos et al., "Multidrug-resistant, extensively drug-resistant and pandrug-resistant bacteria: an international expert proposal for interim standard definitions," *Clinical Microbiology and Infection*, vol. 18, no. 3, pp. 268–281, 2012. <https://doi.org/10.1111/j.1469-0691.2011.03570.x>
- [7] K. Forslund, S. Sunagawa, J. R. Kultima, D. R. Mende, M. Arumugam, A. Typas, and P. Bork, "Country-specific antibiotic use practices impact the human gut resistome," *Genome Research*, vol. 23, no. 7, pp. 1163–1169, 2013. <https://doi.org/10.1101/gr.155465.113>
- [8] B. Kintses et al., "Phylogenetic barriers to horizontal transfer of antimicrobial peptide resistance genes in the human gut microbiota," *Nature Microbiology*, vol. 4, no. 3, pp. 447–458, 2018. <https://doi.org/10.1038/s41564-018-0313-5>
- [9] K. Bush and G. A. Jacoby, "Updated functional classification of β -lactamases," *Antimicrobial Agents and Chemotherapy*, vol. 54, no. 3, pp. 969–976, 2010. <https://doi.org/10.1128/AAC.01009-09>
- [10] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal Statistical Society: Series B*, vol. 57, no. 1, pp. 289–300, 1995. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- [11] V. M. D'Costa, K. M. McGrann, D. W. Hughes, and G. D. Wright, "Sampling the antibiotic resistome," *Science*, vol. 311, no. 5759, pp. 374–377, 2006. <https://doi.org/10.1126/science.1120800>
- [12] I. T. Paulsen et al., "Role of mobile DNA in the evolution of vancomycin-resistant *Enterococcus faecalis*," *Science*, vol. 299, no. 5615, pp. 2071–2074, 2003. <https://doi.org/10.1126/science.1080613>
- [13] D. G. J. Larsson and C. F. Flach, "Antibiotic resistance in the environment," *Nature Reviews Microbiology*, vol. 20, no. 5, pp. 257–269, 2022. <https://doi.org/10.1038/s41579-021-00649-x>