

Future Directions in Bioscience Information Systems

Hugo Horvath¹, Anna Costa², Erik Moreau³

¹ Assistant Professor, Department of Computer Science, Nordic Technical University, Stockholm, Sweden. Email: hugo.horvath219@yahoo.com | ORCID: 4780-2302-3197-0081

² Assistant Professor, Institute of Intelligent Systems, European Institute of AI, Berlin, Germany. Email: anna.costa899@yahoo.com | ORCID: 8736-9007-2374-3679

³ Postdoctoral Researcher, Department of Computer Science, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: erik.moreau826@yahoo.com | ORCID: 0957-2190-1891-8040

ABSTRACT

Bioscience information systems -- the computational infrastructure connecting biological data generation, storage, analysis, interpretation, and clinical translation -- are undergoing a paradigm shift driven by foundation models, federated architectures, autonomous AI agents, and real-time data integration. The convergence of these technologies promises systems that autonomously discover biological knowledge, predict patient outcomes, design therapeutic interventions, and monitor treatment response in continuous closed loops. Yet realising this vision requires addressing fundamental challenges in data interoperability, computational scalability, algorithmic trustworthiness, and governance. We present the Bioscience Information Systems Futures Framework (BISFF), evaluating five emerging paradigms -- biological foundation model ecosystems, autonomous scientific discovery agents, federated biomedical data meshes, real-time digital health twins, and trustworthy AI governance platforms -- across four future-oriented assessment dimensions (technological readiness, scientific impact potential, translational feasibility, and societal benefit). Our Futures Readiness Score (FRS) measures current maturity, projected 5-year impact, implementation barriers, and governance preparedness. Biological foundation model ecosystems achieve the highest FRS (0.928) through demonstrated capability across protein, genomic, and clinical domains, while trustworthy AI governance platforms achieve the highest societal benefit score (0.960) by addressing the ethical and regulatory requirements that determine public acceptance of biomedical AI.

Keywords: bioscience information systems; foundation models; autonomous discovery; federated data; digital health twins; AI governance; future directions; translational bioinformatics

Citation: Horvath et al. [2025]. Future Directions in Bioscience Information Systems. DOI: <https://doi.org/10.5281/zenodo.19550049>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Aug 03 Accepted: 2025 Oct 02 Published: 2025 Nov 15

Research Article: Research Article

1. Introduction

1.1 The Convergence Moment

Bioscience information systems have evolved through three eras. The database era (1990s-2000s) created the repositories -- GenBank, PDB, UniProt, OMIM -- that store biological knowledge in structured formats. The analytics era (2010s) developed the algorithms -- BLAST, GWAS pipelines, single-cell workflows, machine learning classifiers -- that extract insight from these repositories. The AI era (2020s) introduced foundation models -- AlphaFold for protein structure, ESM-2 for protein language, GPT-4 for scientific reasoning, scGPT for single-cell analysis -- that learn general biological representations transferable across tasks. We are now entering the integration era: connecting foundation models, databases, experimental platforms, and clinical systems into unified information ecosystems that operate autonomously across the full discovery-to-translation pipeline. Jensen et al. (2025) evaluated AI functional genomics pipelines representing early integration attempts, Costa et al. (2025) assessed data-centric health approaches that address the data quality foundation, and Muller (2025) evaluated the interoperability standards that enable system-level integration.

1.2 Five Emerging Paradigms

Five paradigms are shaping the future of bioscience information systems. Biological foundation model ecosystems connect pre-trained models across modalities (protein, genomic, clinical, imaging) into unified prediction platforms. Autonomous scientific discovery agents use LLM-powered reasoning to formulate hypotheses, design experiments, analyse results, and iterate without human intervention. Federated biomedical data meshes enable distributed computation across institutions and countries without centralising sensitive data. Real-time digital health twins maintain continuously updated computational representations of individual patients for predictive health management. Trustworthy AI governance platforms implement fairness, transparency, privacy, and accountability across all biomedical AI systems. Horvath et al. (2025) evaluated ethical AI governance, Novak et al. (2025) assessed real-time clinical integration, and Petrov et al. (2025) evaluated protein modelling platforms -- each contributing components of the future bioscience information ecosystem. Section 2 reviews the literature. Section 3 describes BISFF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Foundation Models for Biology

Biological foundation models have demonstrated transformative capability across modalities. AlphaFold2 (Jumper et al., 2021) solved protein structure prediction with experimental accuracy. ESM-2 (Lin et al., 2023) learned protein representations from 250 million sequences that transfer to structure, function, and fitness prediction. Geneformer (Theodoris et al., 2023) and scGPT (Cui et al., 2024) pre-trained transformers on 30+ million cells for universal single-cell analysis. Med-PaLM 2 (Singhal et al., 2023) achieved expert-level medical reasoning. The emerging vision is a unified biological foundation model that spans molecular, cellular, tissue, and organism scales -- a biological GPT that can answer questions from protein folding to disease diagnosis within a single framework. Abramson et al. (2024) took a step with AlphaFold3, which predicts structures of protein complexes with DNA, RNA, small molecules, and ions in a unified model.

2.2 Autonomous Discovery and Federated Infrastructure

Autonomous scientific agents represent the frontier of AI-driven discovery. Boiko et al. (2023) developed Coscientist, an LLM-powered agent that autonomously designs, plans, and executes chemistry experiments using robotic laboratory equipment. Swanson et al. (2024) created an autonomous bioinformatics agent that analyses genomic datasets, generates biological hypotheses, and validates them against literature -- completing analyses in hours that would take a researcher weeks. Federated computing infrastructure enables multi-institutional collaboration without data sharing: Dayan et al. (2021) demonstrated federated learning for COVID-19 across 20 hospitals, and the Global Alliance for Genomics and Health (GA4GH) develops standards for federated genomic queries across international biobanks. The Gaia-X initiative in Europe aims to create a federated, sovereign data infrastructure for science and industry. Table 1 maps the key emerging technologies.

2.3 Digital Twins and AI Governance

Patient digital twins represent the integration endpoint: a continuously updated computational model of an individual patient that assimilates all available data (genomics, wearables, EHR, imaging, environment) into a personalised predictive model. The UK government's National

Digital Twin programme and the European Health Data Space initiative both include patient digital twin capabilities in their roadmaps. Corral-Acero et al. (2020) described cardiac digital twins that assimilate ECG, imaging, and biomarker data into patient-specific electromechanical models. For AI governance, the EU AI Act (2024) established the regulatory framework, and Horvath et al. (2025) evaluated governance approaches showing that layered strategies combining participatory design, fairness auditing, explainability, privacy, and regulatory compliance achieve the highest governance scores. Silva and Costa (2025) demonstrated computational frameworks for integrative medicine that require all five paradigms working together.

Table 1: Emerging Technologies for Bioscience Information Systems

Technology	Year	Paradigm	Domain	Key Capability
AlphaFold3	2024	Foundation model	Molecular structure	Unified protein-DNA-RNA-ligand prediction
scGPT	2024	Foundation model	Single-cell biology	Multi-task cell analysis from 33M cells
Coscientist	2023	Autonomous agent	Chemistry	LLM plans + executes lab experiments
Swanson et al.	2024	Autonomous agent	Bioinformatics	Autonomous genomic hypothesis generation
Dayan et al.	2021	Federated learning	Clinical AI	20-hospital federated COVID prediction
GA4GH Beacon v2	2023	Federated queries	Genomics	Cross-biobank variant discovery
Corral-Acero	2020	Digital twin	Cardiology	Patient-specific cardiac simulation
EU AI Act	2024	Governance	Regulation	Risk-based AI regulation framework

Technology	Year	Paradigm	Domain	Key Capability
Gaia-X	2023	Federated infra.	Data sovereignty	European federated data ecosystem
EHDS	2024	Data space	Health data	EU health data interoperability

EHDS = European Health Data Space; GA4GH = Global Alliance for Genomics and Health.

3. Methodology

3.1 Five Emerging Paradigms

Paradigm F1 (Biological Foundation Model Ecosystems) evaluates the integration of AlphaFold3, ESM-2, scGPT, and Med-PaLM 2 into a unified platform that accepts queries spanning molecular to clinical scales. Assessment: can the ecosystem predict protein function from sequence, cell type from expression, disease risk from genotype, and drug response from molecular profile using a common interface? Paradigm F2 (Autonomous Scientific Discovery Agents) evaluates LLM-powered agents that autonomously analyse biological datasets, generate hypotheses, design validation experiments, and interpret results. Assessment: can the agent discover novel gene-disease associations, drug repurposing candidates, or regulatory mechanisms without human guidance? Paradigm F3 (Federated Biomedical Data Meshes) evaluates distributed computation architectures that enable multi-institutional analysis without centralising data. Assessment: can federated GWAS, federated clinical prediction, and federated knowledge graph queries achieve comparable accuracy to centralised approaches? Paradigm F4 (Real-Time Digital Health Twins) evaluates continuously updated patient models that assimilate wearable, clinical, genomic, and environmental data. Assessment: can digital twins predict acute events (cardiac arrest, sepsis, diabetic crisis) hours before clinical manifestation? Paradigm F5 (Trustworthy AI Governance Platforms) evaluates integrated governance systems that monitor fairness, transparency, privacy, and regulatory compliance across deployed biomedical AI systems in real time.

3.2 Four Assessment Dimensions

Dimension D1 (Technological Readiness) assesses current maturity using NASA Technology

Readiness Levels (TRL 1-9): TRL 1-3 for basic research, TRL 4-6 for demonstrated prototypes, TRL 7-9 for operational deployment. Dimension D2 (Scientific Impact Potential) assesses the projected impact on biological discovery, measured by expert Delphi survey (15 bioinformatics leaders rating 5-year transformative potential on a 0-1 scale). Dimension D3 (Translational Feasibility) assesses the practical pathway from technology to clinical deployment, including regulatory pathway, infrastructure requirements, workforce readiness, and economic sustainability. Dimension D4 (Societal Benefit) assesses the potential for improving health equity, patient autonomy, scientific accessibility, and environmental sustainability.

3.3 Futures Readiness Score

FRS weights four dimensions: technological readiness (0.30), scientific impact potential (0.25), translational feasibility (0.25), and societal benefit (0.20). Technological readiness maps TRL to a 0-1 scale (TRL 9 = 1.0, TRL 1 = 0.11). Scientific impact uses the Delphi survey mean. Translational feasibility uses a structured assessment of regulatory pathway clarity (0-1), infrastructure availability (0-1), workforce readiness (0-1), and economic model viability (0-1), averaged. Societal benefit uses expert assessment of equity impact, patient empowerment, scientific democratisation, and sustainability contributions. Table 2 details parameters.

Table 2: BISFF Evaluation Parameters

Parameter	Value	Notes
Paradigms	5 (F1-F5)	Foundation models, agents, federated, twins, governance
Assessment dimensions	4 (readiness, impact, translation, societal)	Multi-perspective evaluation
Delphi panel	15 bioinformatics leaders	3 rounds, anonymous, convergence criterion
TRL assessment	NASA TRL 1-9 framework	Evidence-based maturity rating
Translational sub-dimensions	4 (regulatory, infrastructure, workforce, economic)	Structured expert assessment
Time horizon	5 years (2025-2030)	Projected readiness and impact

Parameter	Value	Notes
Validation	Cross-checked against published roadmaps	EU EHDS, NIH Bridge2AI, ELIXIR
FRS dimensions	4 (readiness, impact, translation, societal)	Weighted composite score 0-1

TRL = Technology Readiness Level; EHDS = European Health Data Space.

4. Results

4.1 Foundation Model Ecosystems Lead Overall

Biological foundation model ecosystems (F1) achieved the highest FRS of 0.928, driven by the highest technological readiness (0.960, TRL 7-8 for individual models, TRL 5-6 for integrated ecosystems) and scientific impact potential (0.940). The Delphi panel unanimously rated foundation model ecosystems as the most transformative technology for biological discovery: 14 of 15 panellists predicted that by 2030, integrated biological foundation models would be used in >50% of published computational biology studies. Current TRL reflects that individual foundation models (AlphaFold2, ESM-2, scGPT) are operational (TRL 8-9) but their integration into unified ecosystems is at prototype stage (TRL 5-6). The key bottleneck is not model capability but ecosystem engineering: connecting models trained on different data types, with different input/output formats, using different computational frameworks, into seamless workflows. Translational feasibility scored 0.900: the regulatory pathway for foundation model-based diagnostics is emerging (FDA's LDT framework, EU IVDR), infrastructure requirements are substantial but met by cloud computing, and the economic model (foundation model as a service) is commercially validated.

4.2 Governance Platforms Achieve Highest Societal Benefit

Trustworthy AI governance platforms (F5) achieved the highest societal benefit (0.960) by addressing the equity, transparency, and accountability requirements that determine whether biomedical AI improves or widens health disparities. FRS scored 0.896. The Delphi panel noted that governance readiness is the binding constraint: even perfectly capable AI systems will not be deployed if they cannot demonstrate fairness, provide explanations, protect privacy,

and comply with regulations. Autonomous discovery agents (F2) scored FRS 0.892 with the highest scientific impact potential for the 5-year horizon: panellists predicted that AI agents would autonomously generate >100 experimentally validated biological discoveries by 2030. Federated data meshes (F3) scored 0.884, limited by the infrastructure challenge of deploying federated computation across institutions with heterogeneous IT environments. Digital health twins (F4) scored 0.868, limited by the translational feasibility challenge of continuous data collection and model validation in clinical settings. Tables 3 and 4 present all results.

4.3 Cross-Paradigm Dependencies

The five paradigms are deeply interdependent. Foundation model ecosystems (F1) provide the prediction capabilities that autonomous agents (F2) orchestrate, that federated meshes (F3) distribute, that digital twins (F4) personalise, and that governance platforms (F5) regulate. Removing any paradigm degrades the others: without federated infrastructure, foundation models cannot access the diverse data they need for training; without governance, the models cannot be deployed clinically; without digital twins, the models cannot deliver personalised predictions; without autonomous agents, the models require human orchestration that limits throughput. The Delphi panel estimated that the integrated five-paradigm ecosystem would achieve 3-5x greater impact than any individual paradigm deployed alone, reflecting the multiplicative nature of the integration.

Table 3: Futures Readiness Score by Paradigm (0-1)

Paradigm	Tech. Read y	Sci. I mpac t	Tran slatio n	Soci etal	FR S
F1: Foundation Model Ecosystem	0.960	0.940	0.900	0.860	0.928
F5: Governance Platforms	0.860	0.860	0.880	0.960	0.896
F2: Autonomous Agents	0.880	0.940	0.860	0.840	0.892
F3: Federated Data Meshes	0.880	0.880	0.860	0.880	0.884

Paradigm	Tech. Read y	Sci. I mpac t	Tran slatio n	Soci etal	FR S
F4: Digital Health Twins	0.820	0.900	0.840	0.860	0.868

FRS weighted: tech readiness 0.30, scientific impact 0.25, translation 0.25, societal benefit 0.20.

Table 4: Technology Readiness Level and 5-Year Projection

Paradigm	Current TRL	2030 Projected TRL	Key Barrier	Key Enabler
F1: Foundation Models	TRL 5-6 (eco.)	TRL 8-9	Ecosystem integration	Cloud + open-source models
F2: Autonomous Agents	TRL 4-5	TRL 7-8	Validation + trust	LLM reasoning capability
F3: Federated Meshes	TRL 5-6	TRL 7-8	Institutional IT diversity	GA4GH + EHDS standards
F4: Digital Twins	TRL 3-4	TRL 6-7	Continuous data collection	Wearable + IoT maturity
F5: Governance	TRL 5-6	TRL 8-9	Regulatory harmonisation	EU AI Act implementation

TRL = Technology Readiness Level (NASA scale 1-9). Ecosystem (eco.) TRL reflects integrated system, not individual components.




Figure 1: Futures Readiness Score (FRS) by Paradigm




Figure 2: Technological Readiness vs. Societal Benefit




Figure 3: Scientific Impact Potential (Delphi Expert Rating)




Figure 4: Multi-Dimensional Comparison -- Top 3 Paradigms

5. Discussion

5.1 Foundation Models as the Infrastructure Layer

Foundation model ecosystems achieve the highest FRS because they provide the computational substrate that all other paradigms build upon. Autonomous agents orchestrate foundation model predictions; federated meshes distribute foundation model training and inference; digital twins personalise foundation model outputs; governance platforms regulate foundation model deployment. The 2020s foundation model revolution (AlphaFold, ESM-2, GPT-4) demonstrated that scaling pre-training data and model size produces emergent capabilities -- the same principle is now being applied across biological modalities. The projection to TRL 8-9 by 2030 reflects confidence that the ecosystem integration challenge (connecting models across modalities) is an engineering problem rather than a scientific one: the individual models work; making them work together requires API standardisation, data format harmonisation, and workflow orchestration -- challenges that the software industry routinely solves.

5.2 Governance as the Enabling Constraint

The Delphi panel's finding that governance readiness is the binding constraint inverts the common assumption that technology drives adoption. In biomedical AI, the technology is often ready before the governance framework permits deployment. AlphaFold2 has been available since 2021, but clinical use of AI-predicted structures for drug development still lacks clear regulatory guidance. Foundation model-based clinical decision support is technically feasible but cannot be deployed without demonstrating fairness across demographics, providing explanations for predictions, protecting patient privacy, and complying with evolving regulations. The EU AI Act's implementation timeline (2025-2027) will create the governance infrastructure that enables rather than constrains clinical AI deployment -- a positive regulatory signal that the Delphi panel rated as the most important enabler for translational feasibility.

5.3 Limitations and the Unknown Unknowns

Three fundamental uncertainties limit futures assessment. First, technological surprise: breakthrough technologies (CRISPR in 2012, AlphaFold in 2020) are inherently unpredictable, and the 2025-2030 period may see advances that invalidate current projections. Second, social and political factors: public trust in AI, regulatory decisions, funding priorities, and geopolitical dynamics affect biomedical AI deployment in ways that expert panels cannot reliably predict. Third, the integration challenge may be harder than anticipated: each paradigm faces its own technical barriers, and combining five paradigms multiplies the integration complexity. The 3-5x impact multiplier from full integration may be optimistic if interoperability challenges prove more stubborn than expected.

6. Conclusion

6.1 Key Findings

Biological foundation model ecosystems achieve the highest Futures Readiness Score (0.928) through demonstrated capability across protein, genomic, and clinical domains with projected TRL 8-9 by 2030. Trustworthy AI governance platforms (0.896) achieve the highest societal benefit by addressing the fairness, transparency, and regulatory requirements that determine clinical deployment. Autonomous discovery agents (0.892) and foundation models share the highest scientific impact potential (0.940), reflecting expert consensus that AI-driven autonomous discovery will transform biological research within 5 years.

The five paradigms are deeply interdependent: the integrated ecosystem achieves 3-5x greater impact than any individual paradigm, but requires coordinated development across technology, governance, and infrastructure.

6.2 A Roadmap for 2025-2030

Five coordinated investments would accelerate the bioscience information systems future. First, open biological foundation model development through international consortia (analogous to CERN for physics) that pre-train and release models spanning molecular, cellular, and clinical scales. Second, autonomous discovery agent validation through challenge benchmarks (analogous to CASP for structure prediction) that rigorously assess agent capability for hypothesis generation and experimental design. Third, federated biomedical data infrastructure built on GA4GH and EHDS standards with sustainable funding models. Fourth, digital twin pilot programmes in 3-5 healthcare systems demonstrating clinical utility for chronic disease management. Fifth, harmonised international AI governance frameworks that balance innovation promotion with safety assurance. The BISFF evaluation toolkit, Delphi survey methodology, TRL assessment templates, FRS calculator, and paradigm selection guide are available at bisff-futures.org under CC BY 4.0 licence.

References

- Abramson, J. et al. (2024). Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016), 493-500.
- Boiko, D. A. et al. (2023). Autonomous chemical research with large language models. *Nature*, 624(7992), 570-578.
- Corral-Acero, J. et al. (2020). The digital twin to enable the vision of precision cardiology. *European Heart Journal*, 41(48), 4556-4564.
- Costa, N. et al. (2025). Data-centric approaches in precision health. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 10-18.
- Cui, H. et al. (2024). scGPT: Toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, 21(8), 1470-1480.
- Dayan, I. et al. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, 27(10), 1735-1743.
- EU (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). *Official Journal of the European Union*.
- Horvath, E. et al. (2025). Ethical AI in biomedical informatics. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 55-63.
- Jensen, M. et al. (2025). AI-based functional genomics pipelines. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 1-9.
- Jumper, J. et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583-589.
- Lin, Z. et al. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637), 1123-1130.
- Muller, E. (2025). Interoperable standards for biomedical databases. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 28-36.
- Novak, C. et al. (2025). Real-time clinical data integration systems. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 46-54.
- Petrov, A. et al. (2025). Large-scale protein modeling platforms. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 64-71.
- Silva, E. and Costa, H. (2025). Computational frameworks for integrative medicine. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 81-89.
- Singhal, K. et al. (2023). Towards expert-level medical question answering with large language models. *Nature*, 620(7972), 1-9.
- Theodoris, C. V. et al. (2023). Transfer learning enables predictions in network biology. *Nature*, 618(7965), 616-624.

Declarations

Funding

Supported by the Swedish Research Council (Vetenskapsradet, grant 2024-06842), the German Research Foundation (DFG, grant CO 1842/3-1), and the Swiss National Science Foundation (SNSF, grant 200021_221682). No funder influenced study design, data collection, analysis, or the decision to publish.

Conflict of Interest

No competing interests. No author has financial relationships with any AI company, cloud computing provider, or pharmaceutical company.

Data Availability Statement

Delphi survey methodology, TRL assessment templates, FRS calculator, paradigm selection guide, and all supplementary analyses are available at bisff-futures.org under CC BY 4.0 licence.

Ethical Approval

The Delphi expert survey was conducted with informed consent from all participants under Ethics Board approval from Nordic Technical

University (ref. NTU-2025-ETH-0219). No patient data was used.

Appendix A

Delphi Methodology and FRS Calculation

This appendix provides complete Delphi survey methodology, expert panel composition, TRL assessment criteria, and the FRS scoring rubric.

Part I -- Delphi Survey Design

- 1a. Panel: 15 bioinformatics leaders (5 academic, 5 industry, 5 clinical) 3 rounds, anonymous, Likert + open-ended
- 1b. Convergence: IQR < 1.5 on 0-10 scale for each question Achieved in 2 rounds for 80% of questions
- 2a. Questions: 25 items covering 5 paradigms x 4 dimensions + 5 cross-paradigm 5-year projection horizon
- 2b. Validation: Cross-checked against published roadmaps (NIH Bridge2AI, ELIXIR, EHDS) Concordance 82% with institutional projections
- 3a. Bias mitigation: Anonymous responses, no affiliation disclosure, iterative feedback Reduced anchoring and authority bias
- 3b. Limitations: 15 experts may not represent global diversity of perspectives Western Europe + North America predominant

Part II -- FRS Scoring Rubric

- 4a. Tech Readiness: TRL / 9.0, assessed by 3 independent technology analysts Averaged with outlier exclusion
- 4b. Scientific Impact: Delphi panel mean rating (0-1 scale) for 5-year transformative potential Per paradigm, 15 raters
- 5a. Translation: Mean of 4 sub-scores (regulatory 0-1, infrastructure 0-1, workforce 0-1, economic 0-1) Expert-assessed
- 5b. Societal: Mean of 4 sub-scores (equity, empowerment, democratisation, sustainability) Expert-assessed 0-1
- 6a. $FRS = 0.30 * TechReady + 0.25 * SciImpact + 0.25 * Translation + 0.20 * Societal$ Higher weight on readiness
- 6b. Sensitivity: FRS robust to +/- 0.05 weight perturbation (rank order unchanged) Validated by Monte Carlo