CERN IdeaSquare Journal of Experimental Innovation, 2026; 10(2): 107-115
DOI: https://doi.org/10.23726/cij.2026.1830

ORIGINAL ARTICLE

Co-Creation that Leaves Traces in Space Radiobiology

Alessandro Bartoloni[1]*, Lidia Strigari[2]

[1]INFN Sezione di Roma, CERN EP Department, and Sapienza University of Rome, Rome, Italy; [2]IRCCS Azienda Ospedaliero-Universitaria di Bologna, Department of Medical Physics, Bologna, Italy
*Corresponding author: alessandro.bartoloni@cern.ch

ABSTRACT

We report a situated case of interdisciplinary co-creation between astroparticle physics and medical physics/space radiobiology, in which a workable co-creation methodology emerged implicitly from long experience in large scientific collaborations rather than from formal co-creation training. The work is analyzed through three reusable boundary objects that made collaboration testable and transferable: (i) a shared ontology to stabilize meaning and provenance, (ii) dose–effect models to translate physical exposure into health-relevant quantities with explicit validity ranges, and (iii) curated datasets packaged for cross-domain reuse and accountability. We show how these artifacts both enable and constrain agency, surfacing tensions around authorship, invisible labor, and value capture. We propose a minimal evaluation toolkit, provenance-first decision logs, attribution and benefit-sharing ledgers, and lightweight indicators of reuse, semantic convergence, and asymmetry, along with a six-step replication protocol. Finally, we remap the collaboration with a Business Model Canvas to make sustainability assumptions explicit.

Keywords: Co-creation; boundary objects; ontology; dose-effect models; curated datasets; space radiobiology; medical physics; AMS-02; Business Model Canvas.

Received: January 2026. Accepted: August 2026.

INTRODUCTION

Interdisciplinary research is often praised for bringing different forms of expertise together, but collaboration alone does not guarantee that knowledge can move reliably between fields. The practical challenge is to make the assumptions, decisions, and maintenance work of collaboration visible enough to be examined, contested, and reused. Co-creation is often invoked as a positive label. Yet, it can hide asymmetries of power, unresolved questions of authorship and ownership, or inequities in who performs less visible work such as maintenance and curation (von Busch & Palmås, 2023; Drain & Sanders, 2019; Trischler et al., 2019). For interdisciplinary teams, the question is how to make co-creation testable so that collaboration becomes not merely a context for innovation but its engine (Vignoli, 2023).

This paper examines a collaboration at the interface of astroparticle physics and medical physics/space radiobiology, anchored in the INFN Roma Sapienza AMS research group (hereafter, AMS-Roma) and the Medical Physics Division of IRCCS Azienda Ospedaliero-Universitaria di Bologna – Policlinico di Sant'Orsola (hereafter, IRCCS-BO). Throughout this paper, this interdisciplinary collaboration is referred to as the AMS Space Physics and Radiobiology collaboration (AMS-SPRB). The scientific aim is to improve the credibility and reproducibility of radiation health risk assessment for exploratory missions by leveraging high-precision cosmic-ray measurements and translating them into risk-relevant quantities and models (Bartoloni & Strigari, 2021, 2024; Bartoloni et al., 2022a, 2022b, 2022c). The co-creative pathway was not initiated through formal training; it emerged implicitly, almost "unaware", from experience in large scientific collaborations and evidence-driven research practice, and was articulated only later as a replicable approach.

We frame the work as the co-development of three boundary objects: (i) a shared ontology to stabilize meaning and provenance across communities, (ii) a dose-effect modeling stack that translates physical exposure into health-relevant quantities while making assumptions explicit, and (iii) curated datasets that enable reuse, evaluation, and cross-domain accountability. We also propose a minimal evaluation toolkit, a provenance-first decision log, an attribution and benefit-sharing ledger, and lightweight indicators, to make collaboration auditable and replicable. Finally, we remap the collaboration using the Business Model Canvas framework, popularised in Business Model Generation (Osterwalder & Pigneur, 2010), as an accounting tool that surfaces stakeholders, value exchanges, governance assumptions, and the hidden costs of curation and maintenance. The novelty of this work does not lie in inventing new instruments, but in demonstrating how established concepts can be integrated into a coherent empirical methodology for documenting, evaluating, and sustaining interdisciplinary co-creation in Big-Science-to-Health research.

A SITUATED CASE: BIG-SCIENCE-TO-HEALTH TRANSLATION

The Alpha Magnetic Spectrometer (AMS-02) on the International Space Station provides long-term, high-precision measurements of cosmic-ray species over a broad energy range. Space radiobiology and medical physics communities, by contrast, often work with models and datasets organized around biological endpoints, clinical comparators, and radiation-protection frameworks. Bridging these domains therefore requires more than data sharing: it requires a translation framework connecting spectral products and mission scenarios to dosimetric and biologically relevant quantities.

From 2017 onward, the AMS-SPRB collaboration evolved through workshops, joint writing, methodological discussions, and tool development. It did not begin with a fully specified deliverable, but with a shared question: which parts of the radiation health-risk assessment chain are most sensitive to improved cosmic-ray knowledge, and where can cross-domain work generate durable value? Progress emerged iteratively through the production and refinement of artifacts that different communities could legitimately adopt for their own purposes, including shared definitions and mappings, modeling components, and curated datasets. These artifacts became anchors for negotiation, validation, and credit allocation, gradually transforming collaboration from a context for research into a working method.

Methodologically, we treat the AMS-SPRB collaboration as an experimental co-creation process. The empirical material for this article consists primarily of the artifacts themselves, including their successive versions and change histories, together with decision traces reconstructed from manuscript drafts, modeling iterations, data-curation choices, and joint methodological discussions. We do not claim statistical generality. Instead, we aim for analytical transferability by identifying a protocol and a set of lightweight instruments that other Big-Science-to-Health collaborations may adapt. Figure 1 maps the three boundary objects and their relationships to the two disciplinary groups and to the principal outputs of the co-creation process.

Following qualitative process research, we use the term "empirical episode" to denote an analytically bounded segment of the collaboration, reconstructed from versioned artifacts and decision traces, through which a broader co-creation mechanism becomes observable (Langley, 1999; Berends & Deken, 2021). Following Berends and Deken (2021), the analysis primarily adopts a conceptualized composition: the empirical episodes are organized and interpreted through the analytical categories of ontology, dose–effect modeling, and curated datasets, which function as conceptual signposts linking situated events to the broader co-creation framework. The early presentation of the overall boundary-object structure also introduces a limited model-led element, providing readers with a roadmap for relating the three episodes to the evaluation toolkit and six-step protocol developed later in the paper

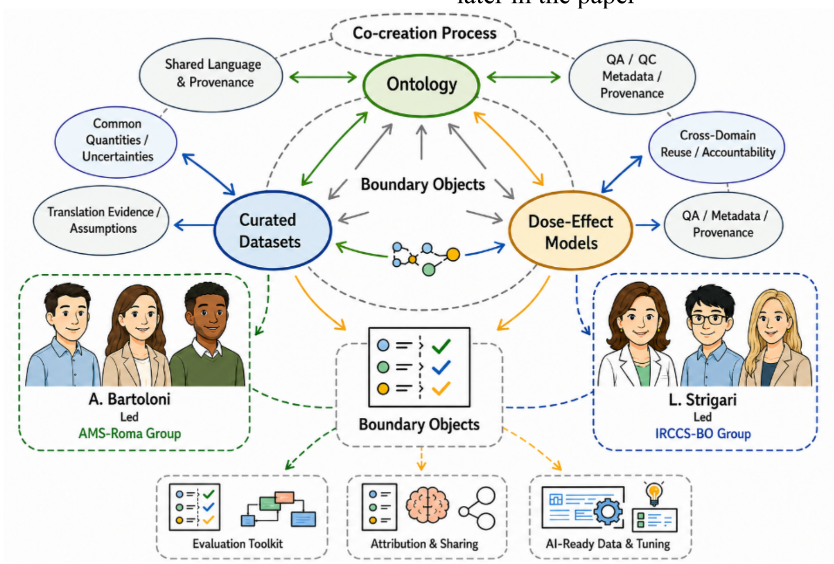

Fig. 1. Boundary objects in the AMS-SPRB co-creation process. Ontology, curated datasets, and dose-effect models mediate collaboration between AMS-Roma and IRCCS-BO. The artifacts support shared language, provenance, translation, reuse, attribution, evaluation, and AI-ready data practices. (Source: Original work of the authors.)

BOUNDARY OBJECT 1: ONTOLOGY AS A MEANING AND PROVENANCE LAYER

A first barrier in interdisciplinary work is language. In space radiobiology, apparently shared terms, such as "dose", "quality factor", "endpoint", "scenario", and "risk", often carry discipline-specific assumptions. The early phase of the collaboration therefore focused on building a shared semantic layer: an ontology, accompanied by a controlled glossary, that encoded the key entities, relations, and minimum metadata required to interpret inputs and outputs. Similar ontology-driven approaches have been proposed for integrating cross-domain knowledge and governance across value networks (Magas et al., 2024).

The ontology served three co-creation functions. First, it stabilized meaning by making definitions explicit and versioned. Second, it supported provenance by linking a claim or model component to its underlying assumptions and data sources. Third, it enabled cross-domain and cross-dataset mapping: cosmic-ray spectral products, transport outputs, dosimetric quantities, and biological endpoints could be represented and linked using a standard set of identifiers.

Empirical episode 1: harmonizing physical quantities for radiobiological use.

An early challenge was that the physical quantities, units, and reporting conventions used in the two domains were not immediately interchangeable. In astroparticle and cosmic-ray physics, AMS-02 results are commonly described in terms of particle species, rigidity or energy intervals, differential fluxes, and observation periods. Although these quantities provide a precise characterization of the measured radiation environment, they cannot be used directly for radiobiological interpretation. Their use in radiobiological analyses required a process of transformation and harmonization, including the alignment of energy-related variables, particle classifications, flux definitions, units, integration intervals, and the assumptions needed to represent a specific mission scenario.

From the medical-physics and radiobiology perspective, the physical description also had to be connected to mission duration, shielding configuration, solar modulation, dose-rate structure, particle mixture, and the biological endpoint under consideration. A related issue concerned the term "dose", which could refer to absorbed dose, dose equivalent, or a health-related quantity or proxy. The difficulty therefore involved not only the meaning of apparently familiar terms, but also the construction of a valid transformation path from AMS-02 measurements to dosimetric and biologically interpretable quantities.

The challenge was not resolved by selecting a single preferred representation. Instead, the collaborators decomposed the translation process into explicit components and associated each physical or dosimetric quantity with its unit, intended use, assumptions, and valid transformation path. The result was a versioned semantic and provenance structure linking measured particle fluxes, rigidity- or energy-based representations, mission conditions, dosimetric quantities, and biological endpoints. This supported the collaboration in two ways: it prevented later modeling steps from silently switching between non-equivalent quantities or incompatible units, and it made data transformation and harmonization visible as scientific work rather than as a merely technical or editorial operation. The modeling consequences of these choices are traceable in Guracho et al. (2023); the associated cell-survival model library and parameter definitions are documented in Supplementary Appendices A and C of that study.

BOUNDARY OBJECT 2: DOSE-EFFECT MODELS AS A TRANSLATION LAYER

The second shared artifact is the modeling layer that translates physical exposure into health-relevant quantities. Dose-effect models served both as scientific tools and as negotiation devices: they required agreement on endpoints, dose metrics, quality factors, validity ranges, and uncertainty treatment. This formalization created a common space in which measured radiation fields could be connected to biological responses without obscuring the assumptions introduced at each step.

A key step was a systematic assessment of dose-effect models used in space radiobiology and related fields (Strigari et al., 2021). The review identified where evidence was sufficiently mature for modeling, where major gaps remained, and which endpoints warranted priority. On this basis, the team developed NTE-DEM as a prototype linking measured exposure inputs, curated biological response functions, and explicit validation choices (Guracho et al., 2023). It was not intended as a complete clinical or mission-risk model; rather, it provided a test environment for examining how physical measurements, biological assumptions, and evidence thresholds could be combined. Technical details are reported in the cited publications. The co-creation function of the modeling layer became particularly visible when the collaborators had to determine what level and type of evidence were sufficient to include a dose–effect relationship in a shared model.

Empirical episode 2: aligning evidence criteria across domains.

Agreement on mathematical implementation did not imply validation agreement. From the astroparticle-physics perspective, exposure inputs had to remain consistent with measured spectra, particle composition, temporal variability, and transport constraints. From the medical-physics and radiobiology perspective, the model also required a biologically meaningful endpoint, an appropriate response function, and evidence that the experimental system was relevant to the intended interpretation.

The initial tendency to treat 'model validation' as a single step was therefore replaced by a layered logic distinguishing: (i) validation of physical exposure inputs; (ii) evidentiary support for the biological response function; (iii) consistency of the integrated model with available ground-based or spaceflight evidence; and (iv) exploratory extrapolations explicitly outside the validated domain. This did not eliminate uncertainty; it located it, assigned responsibility for it, and reduced the risk that an exploratory component would be interpreted as a fully validated health-risk prediction. The resulting chain of evidence is documented in Guracho et al. (2023): Supplementary Appendix A specifies the biological response-model library, Appendix B the proton LET-energy transformation, and Appendix C the parameters used in the hazard functions.

BOUNDARY OBJECT 3: CURATED DATASETS AS AN EVIDENCE AND REUSE LAYER

The third boundary object consists of curated datasets designed for cross-domain reuse. Curated data are often treated as a secondary product of analysis, yet in interdisciplinary projects their structure directly determines which downstream questions can be addressed and under which conditions. Decisions concerning temporal and energy resolution, uncertainty representation, quality controls, dataset structure, documentation, and access therefore constitute substantive scientific choices rather than merely technical preparation.

In Big-Science-to-Health translation, curation has four main requirements. First, datasets must retain sufficient observational and contextual information—including particle species, temporal coverage, orbital and geomagnetic conditions, and relevant solar conditions—to support downstream interpretation. Second, uncertainties and quality indicators must be represented in forms that can be propagated or otherwise incorporated into subsequent analyses. Third, mappings to the shared semantic layer must be maintained so that variables can be connected across datasets without losing their disciplinary meaning. Finally, datasets must be packaged for reuse through explicit versioning, documentation, sharing conditions, and minimal examples showing how the data are intended to be accessed and applied.

Empirical episode 3: balancing measurement fidelity and downstream usability. AMS-02 data products developed for astroparticle analyses preserved levels of temporal, energetic, and species-specific detail that were scientifically appropriate for their original purpose. Downstream radiobiological and dose–effect analyses, however, often required differently structured inputs, including defined temporal summaries, common energy representations, or subsets tailored to particular mission scenarios. The challenge was therefore not simply to transform individual quantities, but to determine how multiple usable representations could be produced without replacing or obscuring the original measurements.

A central decision concerned the level of aggregation. Reducing a time-dependent flux to a mission-average value could facilitate downstream calculations. Still, it could also conceal temporal variability, dose-rate structure, solar modulation, or transient conditions that might be relevant to biological interpretation. Conversely, requiring every downstream user to work directly with the highest-resolution measurements would preserve fidelity but could make reuse unnecessarily difficult and increase the risk of inconsistent processing.

The collaborators addressed this tension by retaining the higher-resolution source data while providing explicitly derived dataset views for defined modeling purposes. Each derived representation was linked to its source version and accompanied by information on its intended use, aggregation procedure, uncertainty treatment, temporal and energy coverage, and limitations. Exclusions and quality-control decisions remained visible rather than being absorbed into an apparently complete dataset. The resulting data resource was therefore not simply a cleaned version of the original measurements. It constituted a layered and versioned evidence object that balanced measurement fidelity with usability for downstream dosimetric and biological analyses.

CROSS-CUTTING IMPLICATIONS: MAKING COLLABORATIVE LABOR VISIBLE

Across the three boundary objects, the work enabling cross-domain interpretation and reuse was distributed unevenly and was not always visible in conventional research outputs. In this context, value has three complementary meanings: formal scientific recognition, including authorship, acknowledgments, and citations; influence over scientific decisions and interpretation; and the allocation of time, funding, and maintenance resources.

The relevant labor extends beyond conventional data analysis. It includes semantic mediation, glossary and ontology maintenance, literature screening, dataset curation, uncertainty annotation, documentation, meeting facilitation, software maintenance, and validation-workflow design. These contributions rarely appear as headline scientific results, yet they determine whether shared artifacts can be interpreted, reused, maintained, and trusted. A related example is the replica of the Alpha Magnetic Spectrometer Payload Operations Control Room, for which scientific translation, educational design, material preparation, and stakeholder engagement were integral to developing a reusable outreach infrastructure (Bartoloni et al., 2024). The proposed ABSL is intended to make such contributions visible prospectively rather than reconstructing them only after project completion.

AI may support selected components of this work, including candidate metadata extraction, ontology mapping, anomaly detection across curated datasets, and structured documentation. These applications may accelerate comparison and consistency checks, but they do not transfer scientific authority to the system. Human experts retain responsibility for definitions, inclusion and exclusion criteria, model endpoints, validity ranges, uncertainty treatment, and quality thresholds. AI-generated outputs should therefore be treated as provisional propositions requiring review against provenance records and domain expertise, reducing the risk of "false agreement", in which fluent synthesis conceals unresolved disciplinary differences.

SYNTHESIS ACROSS THE THREE BOUNDARY OBJECTS

Taken together, the three empirical episodes show that the shared artifacts did not arise as finished products. They emerged through repeated attempts to address different types of incompatibility: how physical quantities could be interpreted across domains, what evidence was sufficient to support a dose–effect relationship, and how measurement products could be structured for reuse beyond their original disciplinary setting.

In each case, the outcome was not only a scientific decision but also a persistent and inspectable trace: a definition, mapping, validation criterion, metadata field, dataset view, transformation rule, uncertainty statement, or documented exclusion. These traces form the empirical basis for the evaluation toolkit proposed in the following section.

Table 1. Empirical episodes underlying the co-creation framework.

Episode	Cross-domain challenge → negotiated trace → downstream effect
1) Harmonizing physical quantities for radiobiological use	Non-interchangeable physical quantities, units, and reporting conventions → versioned semantic and provenance structure linking measured particle fluxes, rigidity- or energy-based representations, mission conditions, dosimetric quantities, and biological endpoints → prevented silent shifts between non-equivalent quantities or incompatible units.
2) Model validation: Aligning evidence criteria across domains	Mathematical and physical consistency did not necessarily imply biological relevance or adequate validation → layered validation logic distinguishing physical-input validation, evidentiary support for biological response functions, integrated-model consistency, and explicitly flagged extrapolations → reduced the risk that exploratory components would be interpreted as validated health-risk predictions.
3 Balancing measurement fidelity and downstream usability	Preservation of high-resolution measurements versus the need for differently structured downstream inputs → layered and versioned dataset views linked to source data, with documented aggregation procedures, uncertainty treatment, coverage, limitations, exclusions, and intended uses → enabled cross-domain reuse without obscuring measurement fidelity or analytical limitations.

MINIMAL EVALUATION TOOLKIT AND SIX-STEP PROTOCOL

To make co-creation rigorous, interdisciplinary teams need instruments that are lightweight enough to adopt, yet strong enough to support evaluation, reuse, and fair attribution. In our case, many of the practices described below were first enacted implicitly, emerging from habits learned in extensive scientific collaborations and evidence-driven research, before being articulated explicitly. We therefore present them as a next step: a deliberate move from "working co-creatively" to leaving usable traces that others can inspect, replicate, and build upon. Importantly, these instruments should not be read as proprietary inventions. They draw on established ideas in participatory design, collaboration evaluation, and provenance-aware research practice. Still, they are here remapped and operationalized around our three boundary objects (ontology, dose–effect models, curated datasets) to fit a Big-Science ↔ Health translation context.

Table 2 summarises the proposed minimal toolkit: (i) a provenance-first decision log (PDL) that links key choices to specific ontology terms, model components, and dataset versions; (ii) an attribution and benefit-sharing ledger (ABSL) that makes visible the contributions typically under-credited (curation, mapping, facilitation) and defines reuse terms; and (iii) a small set of minimal co-creation indicators (MCI), reuse across domains, semantic convergence, and an asymmetry flag, that can be computed periodically to detect drift, persistent mismatch, or concentration of agency. These elements are intentionally minimal: the goal is not bureaucratic overhead, but preserving the traceability required for replication and equitable credit allocation, core values of experimentation-driven innovation.

Building on this toolkit, we propose a six-step replicable protocol that other teams can adapt as a practical guide to "traceable co-creation."

- First, frame the shared problem with dual vocabularies and identify where translation is required.
- Second, define the boundary objects early and agree on what "reuse" will mean in practice.
- Third, establish governance from the outset, including decision rights, attribution rules, and benefit-sharing principles.
- Fourth, build and version the ontology as a living contract for meaning and provenance.
- Fifth, develop translation models with explicit validity ranges and validation targets.
- Sixth, curate datasets with minimal metadata, quality checks, and reference workflows, and evaluate progress periodically using the minimal indicators.

The protocol does not prescribe scientific content; it prescribes how to organize collaboration so that decisions remain legible, artifacts remain reusable, and co-creation remains accountable across disciplinary boundaries.

An illustrative PDL entry would record the decision to describe a mission scenario through duration, shielding configuration, solar modulation, and activity profile, while documenting the rejected alternative of elapsed mission time alone. The corresponding ABSL entry would recognize semantic mediation, ontology development, data preparation, curation, and validation work. The MCI would then monitor cross-domain reuse, semantic stability across iterations, and whether decision authority became concentrated within one discipline.

Table 2. Minimal evaluation toolkit for traceable co-creation.

Instrument	What it records/measures	How it is used	What it detects
Provenance-first Decision Log (PDL)	Key choices, alternatives, evidence, scope, version, owner; links to ontology terms, model components, and dataset versions	Updated at each significant decision; referenced in manuscripts, code releases, and dataset packages	Silent assumption drift; unexamined trade-offs; loss of rationale across iterations
Attribution & Benefit-Sharing Ledger (ABSL)	Contributions including facilitation, curation, mapping, modeling, credit rules, and reuse terms	Reviewed at milestones (submission, release); used to allocate authorship/acknowledgements and define reuse conditions	Invisible labor; value capture; inequitable credit distribution
Minimal Co-creation Indicators (MCI)	Reuse across domains; semantic convergence; asymmetry flag	Computed periodically from artifact usage and decision traces; used to compare iterations or teams	Low reusability; persistent semantic mismatch; concentration of agency

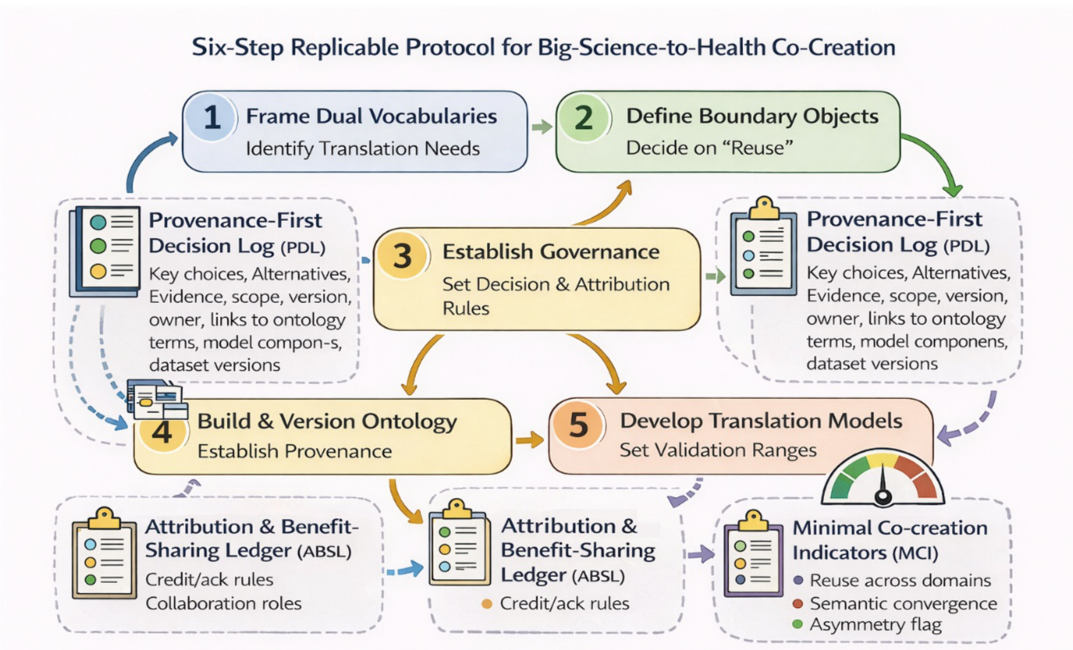

Fig. 2. Six-step protocol for traceable Big-Science-to-Health co-creation. PDL, ABSL, and MCI are embedded across the iterative workflow to support provenance, attribution, and periodic evaluation. (Source: Original work of the authors.).

REMAPPING CO-CREATION THROUGH THE BUSINESS MODEL CANVAS: RESEARCH AS A VALUE SYSTEM

As a final reflective exercise, we remap the collaboration using the Business Model Canvas (BMC) framework, popularised in *Business Model Generation* (Osterwalder & Pigneur, 2010). In this context, the Canvas is not used as a commercialization tool. Instead, it functions as an analytical and accounting framework that makes explicit the structure of value creation, exchange, governance, and sustainability within a research-driven collaboration. In other words, the "business" described here is research itself: the production, maintenance, recognition, and reuse of scientific knowledge across disciplinary boundaries. This perspective is particularly appropriate because the Business Model Canvas itself emerged through a large-scale co-creation process involving an international community of practitioners and researchers. We therefore use a co-created framework to reflect upon another co-created process.

The BMC adds an organizational perspective that complements the boundary-object analysis. Whereas boundary objects explain how meaning, evidence, and models travel across disciplinary communities, the Canvas reveals who sustains that circulation, where value accumulates, which resources remain critical or fragile, and which costs may otherwise be externalized as invisible academic labor. It therefore provides a complementary lens for examining governance, sustainability, responsibility, and long-term value creation.
Figure 3 summarises this remapping.

In a research context, **Customer Segments** correspond to communities of use and responsibility rather than buyers, including space agencies and radiation-risk groups, radiobiology researchers, medical-physics and radiotherapy communities, model developers, tool builders, and standards bodies.

The **Value Proposition** is a translation framework connecting high-precision cosmic-ray measurements to health-relevant quantities through a shared ontology, dose–effect models, curated datasets, and evaluation tools.

Channels include peer-reviewed publications, data repositories, workshops, conferences, and reference workflows through which research outputs circulate.

Customer Relationships are based on long-term collaboration, co-authorship, trust, mentorship, shared stewardship, and cross-training.

Key Resources include AMS data and expertise, the ontology framework, dose–effect model codebases, curated datasets, computational infrastructure, and the human expertise required to connect astroparticle physics, medical physics, radiobiology, and data stewardship.

Key Activities include ontology and dataset curation, dose–effect modeling, cross-domain translation, validation, documentation, governance, and maintenance.

Key Partners include the INFN/AMS Collaboration, Sapienza University of Rome, CERN, the IRCCS Azienda Ospedaliero-Universitaria di Bologna–Policlinico di Sant'Orsola, medical-physics and clinical research groups, space-health programs, and external developers of models and transport codes.

In this research-oriented Canvas, **Revenue Streams** are more appropriately understood as sustainability and recognition streams. They include competitive research funding, institutional and in-kind contributions, citations, integration into external tools and models, and uptake of reference workflows. A further challenge concerns the sustainability of genuinely interdisciplinary activities. Research funding and reporting structures are often organized around disciplinary programs, whereas co-creation may require resources, expertise, and outputs that span several domains simultaneously. Activities such as semantic integration, shared data stewardship, and cross-domain validation may therefore generate substantial scientific value without fitting neatly within the objectives, cost categories, or reporting logic of a single funding stream. Making these contributions visible helps explain why an organizational and sustainability perspective is needed alongside the analysis of boundary objects.

The **Cost Structure** makes visible the human and infrastructural costs of co-creation, including personnel, data storage and computation, software and tool maintenance, dataset curation, semantic mediation, documentation, meetings, workshops, dissemination, and coordination across institutions.

Used in this way, the BMC does not replace the boundary-object framework, but complements it. The Canvas reveals the organizational and sustainability conditions of collaboration, whereas the boundary-object perspective explains how technical and epistemic translation occurs. It should therefore be understood as a reflective follow-up to the collaboration—a structured way of examining, after the fact, how value, labor, responsibility, resources, and sustainability were distributed—rather than as a tool for commercializing the research.

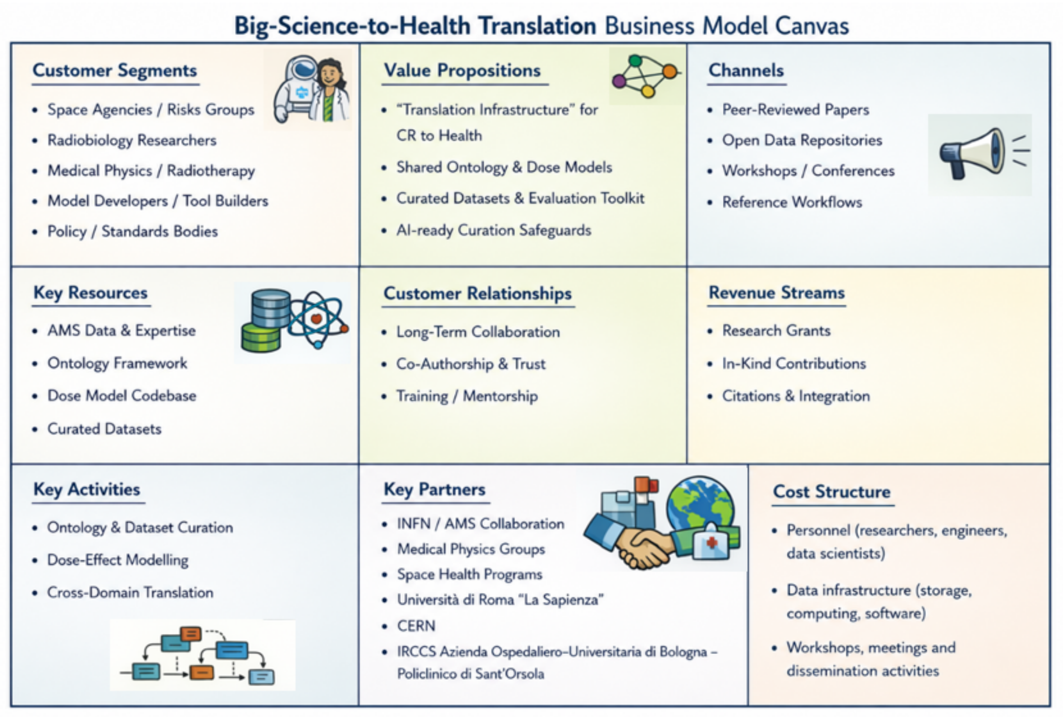

Fig. 3. Big-Science-to-Health Translation Business Model Canvas. The Canvas treats research as a value system and makes stakeholders, value flows, sustainability mechanisms, and hidden costs explicit. (Source: Original work of the authors. The Canvas content was adapted from Osterwalder and Pigneur (2010) to the research case).

CONCLUSIONS AND NEXT STEPS

Co-creation becomes an engine of innovation when it produces artifacts that can travel across disciplinary boundaries while remaining accountable: their meanings, assumptions, transformations, and conditions of reuse can be inspected, challenged, and revised, and in the AMS-SPRB collaboration, ontology development, dose–effect modeling, and curated datasets acted as complementary boundary objects. Together, they enabled translation from cosmic-ray measurements to biologically relevant interpretations while also making visible tensions concerning disciplinary authority, evidentiary standards, attribution, and value. The proposed minimal evaluation toolkit and six-step protocol translate lessons that initially emerged through practice into an explicit framework for traceable co-creation. By integrating a provenance-first decision log, an attribution and benefit-sharing ledger, and minimal co-creation indicators across the development of shared artifacts, interdisciplinary teams can improve reproducibility, reuse, and recognition without imposing excessive procedural burden or constraining scientific experimentation.

The next step is prospective validation. The toolkit should first be implemented from the outset in new Big-Science-to-Health activities, rather than reconstructed retrospectively, so that its feasibility, workload, and ability to preserve decision traces can be assessed over time. It should then be tested in other interdisciplinary settings to determine which elements are transferable and which require adaptation to different scientific, organizational, and governance conditions. Comparative applications could also support refinement of the minimal indicators, particularly their capacity to detect semantic drift, limited cross-domain reuse, invisible labor, and concentration of decision authority.

Future work should therefore focus on three priorities: evaluating the protocol longitudinally across multiple projects; developing transparent and practically measurable definitions for the proposed indicators; and examining whether the recorded traces improve artifact reuse, reproducibility, and equitable attribution. AI may assist with metadata extraction, ontology alignment, documentation, and consistency checks, but its contribution should also be evaluated prospectively. Such systems should remain subordinate to documented human judgment and should preserve, rather than conceal, legitimate disciplinary disagreement.

The broader objective is not to standardize the scientific content of interdisciplinary collaboration. It is to establish a sufficiently robust and adaptable infrastructure through which decisions remain legible, shared artifacts remain reusable, and responsibility for scientific and organizational choices remains visible. In this sense, the proposed framework should be considered not a finished model, but a testable foundation for a wider research agenda on accountable and traceable co-creation.

ACKNOWLEDGEMENTS

This article builds on collaborative work carried out within the INFN Roma Sapienza AMS research group and its medical physics and space radiobiology collaborations. We acknowledge the broader AMS Collaboration for enabling the high-quality cosmic-ray measurements that motivate the translation agenda, and we thank the colleagues and workshop participants who contributed to cross-domain discussions and the development of early prototypes.

CONFLICTS OF INTEREST

None to declare.

REFERENCES

Bartoloni, A., Della Gala, G., Guracho, A. N., Paolani, G., Santoro, M., Strigari, L., & Strolin, S. (2022a). Astroparticle experiments to improve the radiation health risk assessment for humans in space missions. *Proceedings of the International Astronautical Congress (IAC2022).*

Bartoloni, A., Della Gala, G., Guracho, A. N., Paolani, G., Santoro, M., Strigari, L., & Strolin, S. (2022b). Dose-Effects Models for Space Radiobiology: an overview on Dose-Effect Relationship. *Proceedings of the International Astronautical Congress (IAC2022).*

Bartoloni, A., Paolani, G., Santoro, M., Strigari, L., Strolin, S., Guracho, A. N., Della Gala, G., (2022c). High-energy physics astroparticle experiments to improve radiation health risk assessment in space missions. *PoS (EPS-HEP2021)*. https://doi.org/10.22323/1.398.0106

Bartoloni, A., Peroni, M., Strigari, L., & Guracho, A. N. (2024). "So far, so close": Communicating science through a replica of the Alpha Magnetic Spectrometer Payload Operation Control Room. Proceedings of the European Physical Society *Conference on High Energy Physics (EPS-HEP2023), PoS(EPS-HEP2023)*, 648. https://doi.org/10.22323/1.449.0648

Bartoloni, A., & Strigari, L. (2021). Can high-energy particle detectors be used for improving risk models in space radiobiology? Proceedings of the Global Exploration Forum (GLEX2021).

Bartoloni, A., & Strigari, L. (2025). Space radiobiology: Synergies between astroparticle and medical physics. World Scientific.

Berends, H., & Deken, F. (2021). Composing qualitative process research. *Strategic Organization*, 19(1), 134–146.

Drain, A., & Sanders, E. B. N. (2019). A collaboration system model for planning and evaluating participatory design projects. *International Journal of Design*, 13(3), 39–52.

Guracho, A. N., Strigari, L., Della Gala, G., Paolani, G., Santoro, M., Strolin, S., & Bartoloni, A. (2023). Space radiation-induced bystander effect in estimating the carcinogenic risk due to galactic cosmic rays. *Journal of Mechanics in Medicine and Biology*, 23(6), 2340023. https://doi.org/10.1142/S0219519423400237

Langley, A. (1999). Strategies for theorizing from process data. *Academy of Management Review*, 24(4), 691–710.

Magas, M., Kiritsis, D., Poveda-Villalón, M., Yang, L., Björling, S.-E., & Rudenå, A. (2024). Ecosystem integration: The use of ontologies in integrating knowledge across manufacturing value networks. *Frontiers in Manufacturing Technology*, 4, 1331197. https://doi.org/10.3389/fmtec.2024.1331197

Osterwalder, A., & Pigneur, Y. (2010). *Business model generation: A handbook for visionaries, game changers, and challengers*. John Wiley & Sons.

Strigari, L., Strolin, S., Morganti, A. G., & Bartoloni, A. (2021). Dose-effects models for space radiobiology: An overview on dose-effect relationships. *Frontiers in Public Health*, 9, 733337. https://doi.org/10.3389/fpubh.2021.733337

Trischler, J., Dietrich, T., & Rundle-Thiele, S. (2019). Co-design: From expert- to user-driven ideas in public service design. *Public Management Review*, 21(11), 1595–1619.

Vignoli, M. (2023). New research methodologies in innovation: A shift toward experimentation. *CERN IdeaSquare Journal of Experimental Innovation*, 7(2), 1–3. https://doi.org/10.23726/cij.2023.1482

von Busch, O., & Palmås, K. (2023). *The corruption of co-design: Political and social conflicts in participatory design thinking*. Routledge.