

Interoperability and the Democratization of Care

Industry scale, equitable participation and the economics of connected services

RESEARCH LENS	SCOPE
International / Regional / Cross-specialty	Industry scale, public benefit, provider universes and sustainable interoperability
Central question	How can this field advance the democratization of care?
Evidence cutoff	October 2, 2026

Democratization of care means expanding people's practical ability to obtain appropriate, affordable care, exercise informed choice and participate in decisions that affect them.

An AI-assisted academic working paper. General subject-matter research; no original participant data or independent peer review. Empirical evidence, mathematical results, assumptions and proposed evaluations are distinguished in the text.

Abstract & principal findings

Public-interest focus. This paper examines interoperability as a means of advancing democratization of care: wider practical access, affordability, informed choice and equitable participation. Claims of benefit require evidence about who gains and who remains excluded.

Research question. How large is the economic opportunity for interoperability across healthcare and adjacent care services, and which portion could plausibly support a sustainable interoperability service? This paper separates industry expenditure, fragmented workflows, identifiable purchasing organizations and attainable platform revenue. Its central finding is methodological: a large care economy establishes relevance, but does not establish a software total addressable market.

Approach. A targeted narrative review combines official expenditure statistics, US hospital interoperability evidence, business-statistics definitions and a critical audit of the supplied provider-universe document. The analysis then develops original account-based sizing equations, scenario sensitivity, workflow selection criteria and a prospective validation program. It does not conduct a new global census or estimate a statistically representative worldwide adoption rate.

Evidence. WHO reports approximately US\$9.8 trillion in global health spending for 2022. CMS reports US\$5.3 trillion in US health expenditure for 2024. These are different-year, overlapping geographic aggregates and must not be added. ASTP/ONC describes substantially different rates for sometimes-or-routinely versus routinely exchanging information across all four measured domains in US hospitals.^{[1][2][3]}

Interpretation. The supplied document's hundreds-of-millions provider universe is not a defensible buyer count. Its categories combine workers, practices and facilities, while its overlap adjustments and interoperability percentages are unsupported. Those ranges should be retained as internal hypotheses until independently reconstructed.^[6]

Contribution. The proposed research program replaces a headline-driven market narrative with a measurable progression: define a buyer, verify an unmet workflow, establish legal and technical access, prove an economic benefit, and observe willingness to pay. Interoperability is valuable where it completes a useful transaction under accountable conditions. Connectivity alone is an intermediate output.

Principal conclusion: Interoperability can advance democratization of care when it reduces exclusion, makes coordination affordable and supports locally sustainable services. Market size should be evaluated separately from public benefit.

Reading map

Sections 1-3 establish units and audit the original evidence. Sections 4-7 develop market and operating models. Sections 8-10 define regional, specialty and evaluation priorities. Appendices provide a reproducible data specification and a claim-release checklist.

Research method & evidence appraisal

Market evidence and analytical validity

The review used purposive source selection to answer three separate questions: the scale of health activity, the extent of a measured interoperability gap and the construction of a commercial buyer denominator. Official statistics were prioritized for expenditure and business-unit definitions. The supplied planning PDF was treated as an object of audit rather than independent corroboration. No new provider census was performed.

Evidence class	Extraction focus	Interpretation limit
Expenditure statistics	Geography, year, currency and sector boundary	Context only; no platform capture assumed
Hospital exchange survey	Population and frequency definition	No extrapolation to all care sectors
Business statistics	Establishment versus nonemployer unit	Method reference; new counts not extracted
Internal planning source	Traceability and overlap adjustment	Unsupported ranges remain hypotheses

A claim was retained as an empirical baseline only when its source, unit and observation period were identifiable. Scenario inputs were kept separate from those baselines. Arithmetic was checked directly, while external datasets were not independently reconstructed. The analysis uses dimensional consistency as a validity test: people, organizations, transactions and dollars cannot be summed or exchanged without an explicit mapping.

The main sources of uncertainty are definition mismatch, selection bias, geographic transfer and unknown commercial behavior. The sensitivity analysis is deterministic and does not assign probabilities to scenarios. It should not be read as a statistical confidence interval. A future empirical model should publish its extraction code, deduplication rules and uncertainty treatment.

Transparency statement

This is an AI-assisted, targeted narrative synthesis. Source selection was purposive rather than exhaustive; no registered systematic-review protocol is claimed. The reference list identifies the material used, including access limitations. No original participant data were collected. Proposed studies, capabilities and economic scenarios are labeled and have not been independently validated.

1. Define the object being measured

Interoperability has several objects: a technical connection, a semantically meaningful data exchange, a legally permissible disclosure and a completed organizational workflow. Counting one does not establish another. A system may expose an API while omitting appointment availability; a practice may exchange clinical documents yet still fax referral status; a directory may list an organization whose scheduling endpoint is unavailable.

For this paper, **operational interoperability** means that authorized participants can exchange the minimum necessary information and complete a defined cross-organizational task with interpretable results, accountable ownership and recovery from failure. This is a proposed analytical definition, not a claim that a single standard certifies the whole process. It deliberately makes the outcome, rather than the interface, the unit of value.

The denominator must match the business model. Per-clinician software may use licensed professionals after removing inactive and duplicate records. An enterprise platform may sell to a parent organization operating many facilities. A transaction service may price successful events while negotiating with a comparatively small number of buyers. These universes cannot be exchanged merely because all are called providers.

Unit	Appropriate use	Main counting risk
Person or professional	Workforce, seats, credentials	Multiple licenses and practice locations
Establishment	Local service availability	Several establishments share one buyer
Legal organization	Contracting and liability	Parent-child consolidation obscures procurement
Buying account	Sales and revenue scenarios	Budget owner may differ from operator
Endpoint or transaction	Integration and usage	Many endpoints per buyer; retries inflate events

Business demography should therefore be recorded as a relational dataset, with entities connected to establishments, people and endpoints. The proposed primary key is a verified business identifier within a jurisdiction, supplemented by explicit parent-child relationships. Fuzzy name matching can nominate duplicates; it should not silently collapse distinct organizations. Every merge needs a confidence score and a reversible record.

The resulting account universe will be much smaller than a broad population of care workers in some sectors and more fragmented in others. That difference is not a weakness of the market. It determines the actual distribution of sales effort, onboarding cost, purchasing power and operational responsibility.

2. Industry scale without false precision

Health expenditure establishes the scale of activity whose coordination might improve. It includes labor, facilities, medicines, services and other spending; it is not the revenue pool available to an interoperability platform. Even a perfectly measured expenditure total would not tell us which organizations can buy software, how much they will pay or which workflows a new entrant can access.

Baseline	Reported value	Interpretation boundary
Global health expenditure, 2022	Approximately US\$9.8 trillion ^[1]	Context for health activity; not a 2026 TAM
US health expenditure, 2024	US\$5.3 trillion ^[2]	US subset of the world; different observation year
US GDP share, 2024	18.0% ^[2]	Economic intensity, not platform capture potential
US out-of-pocket spending, 2024	US\$556.6 billion ^[2]	Affordability context; not available marketplace GMV

These aggregates should appear in separate panels in external materials, with geography and observation year attached to the value itself. A footnote alone is too easy to detach when a slide or webpage is reused. Forecasts, if introduced later, require a different visual treatment from observed statistics and an explicit forecast horizon.

Adjacent care services present another boundary problem. Childcare, pet care, personal wellness, beauty and fitness can share discovery, booking and payment infrastructure with healthcare. Their expenditure and business counts nevertheless belong to different statistical classifications. Calling all of them care is a legitimate product strategy; treating one healthcare aggregate as their combined economic size is not.

The proposed sizing architecture consequently maintains two linked ledgers. The **activity ledger** records expenditure, service volumes and unmet need. The **commercial ledger** records eligible buyers, reachable accounts, procurement constraints, willingness to pay and retention. A documented mapping connects them when a specific workflow supports a payment. Without that mapping, expenditure remains context.

A conservative publication rule

Publish only a market figure whose unit, geography, year, source and calculation are reproducible. If no credible global buyer count exists for the intended product, publish a bounded regional analysis and explain the missing evidence. A transparent incomplete estimate is more useful for decision-making than a precise-looking global total assembled from incompatible units.

3. Audit of the supplied opportunity document

The supplied seven-page planning document is a useful inventory of sectors but an insufficient empirical basis for a market-size claim. It reports roughly 392-633 million or more providers/facilities in one summary; the 19 visible category endpoints in the companion table actually sum to 392-632 million. This one-million upper-end discrepancy should be corrected before reuse. Additional gross, unique and digitally fragmented ranges are not reproducible from a documented linkage method. Its categories include medical, personal, family and animal care. Supporting sources, observation years and reproducible overlap adjustments are not supplied.^[6]

Original element	Audit finding	Required disposition
Global provider/facility ranges	Workers, organizations and facilities are mixed	Do not present as verified buyers
Unique universe after overlap	No entity-level linkage or adjustment method	Rebuild from defined statistical units
Online / digital / interoperable shares	No sample, instrument or geographic weighting	Treat as hypotheses for field research
Opportunity as 100% minus interoperability	A technical gap is not purchase intent	Add budget, access and value filters
US buyer/business ranges	No traceable extraction or category reconciliation	Replace with documented regional data

The most important correction is not choosing a slightly smaller total. It is changing the model. A hospital, a physician employed by that hospital and the hospital's outpatient site may represent three listed providers but one purchasing decision. Conversely, a franchise network may share branding while locations purchase independently. The direction of error cannot be resolved by applying a universal discount.

There is also a conceptual overlap between “digitally strong” and “interoperable.” A business may have polished consumer booking and poor clinical exchange, or excellent record exchange and a weak online presence. These dimensions should be measured separately. A single maturity percentage conceals the actual product need and can reverse the apparent priority of a segment.

Preserve useful structure; retire unsupported totals


The sector list should become a research taxonomy. For each category, record an authoritative source, unit, geography, year, ownership structure, likely budget holder and workflow hypothesis. Do not automatically sum the rows. A sum becomes meaningful only after every row is expressed in a compatible unit and overlapping membership is resolved.

The original document should remain archived with a clear “unverified planning estimates” label. Subsequent versions should identify which claims were replaced and why. This preserves institutional learning while preventing old ranges from reappearing as factual statements in marketing, investor materials or academic-looking reports.

4. What measured interoperability tells us

ASTP/ONC's 2023 hospital evidence distinguishes occasional capability from routine execution. Seventy percent of US nonfederal acute care hospitals reported sometimes or routinely engaging in all four measured exchange domains; 43% did so routinely. The remaining 27 percentage points represent the sometimes category, leaving 30% outside the all-four grouping.^[3]

US NONFEDERAL ACUTE CARE HOSPITALS | 2023



Find + send + receive + integrate. National hospital results, not a global rate.

Source: ASTP/ONC [3]. Categories are mutually exclusive; total = 100%.

These results support a nuanced opportunity thesis. The industry is neither wholly disconnected nor fully integrated. More advanced organizations may still need reliable orchestration, reconciliation and exception handling. Less connected organizations may first need usable records, staff training and basic workflow digitization. Selling both groups the same product obscures their different readiness and implementation burdens.

The four-domain measure is also not an end-to-end measure of every patient journey. A hospital's capacity to receive and integrate data does not reveal whether a referral was accepted, a patient could afford the visit or the receiving service had capacity. Those outcomes need separate instruments. An interoperability research program should avoid presenting a technical maturity score as a direct access-to-care outcome.

Proposed maturity instrument

Dimension	Observation to collect	Failure concealed by a binary label
Connection	Supported transport and availability	Endpoint exists but is unreliable
Meaning	Codes, units and provenance preserved	Syntactically valid, clinically ambiguous data
Authority	Permitted actor, purpose and recipient	Accessible data used beyond allowed scope
Completion	Receipt, disposition and follow-up	Message delivered but task abandoned

The survey is a US hospital baseline. It does not validate the supplied document's global percentages for dental, beauty, childcare or home care. Applying it to those sectors would create apparent rigor without a valid denominator. Instead, it can inform the design of a sector-specific survey that asks about concrete tasks and records the evidence behind each response.

Self-report should be complemented by a small transaction audit. Ask a participating organization to demonstrate a recent completed workflow, including failures and handoffs. The combination of declared capability and observed execution is a stronger basis for product design than either measure alone.

5. A bottom-up market model

The proposed commercial model begins with eligible buying accounts rather than industry spending. Let N denote independently deduplicated accounts in a defined segment; r the reachable fraction; e the fraction that can be technically and legally enabled; c the paid conversion among enabled reachable accounts; and p annual recurring revenue per paid account. Then **modeled ARR** = $N \times r \times e \times c \times p$. Each parameter must have a clear conditional denominator.

Illustrative parameter	Assumption	Evidence needed before external use
N: eligible accounts	150,000	Reconciled business register and exclusion rules
r: reachable share	40%	Measured channel access and qualified contact rate
e: enablement share	50%	Integration, contract and operational screening
c: paid conversion	20%	Observed conversion within the enabled cohort
p: annual revenue	US\$1,200	Signed pricing and realized revenue definition

Under these deliberately hypothetical inputs, 60,000 accounts are reachable, 30,000 are enabled, and 6,000 become paid accounts. Modeled annual recurring revenue is US\$7.2 million. This is a worked calculation, not a measured TAM, company forecast or assessment of an existing customer base.

The ordering matters. If a conversion estimate already includes technical failures, multiplying by a second enablement factor would double-discount. If an annual price is actually a one-time membership payment, classifying it as ARR would overstate recurring revenue. The model's dictionary must define both the event and the time period before numbers are entered.

Separate three economic claims

Addressable accounts are eligible buyers with a relevant need. **Serviceable accounts** can be reached and supported within the product's geography and operating constraints. **Attainable revenue** additionally requires adoption, payment and retention over time. These layers can be reported together, but not conflated into a single large headline.

A mature model adds time: cohort acquisition, onboarding delay, churn, expansion and implementation capacity. A company cannot serve all theoretically available accounts immediately. A quarterly account-flow model often reveals an operational bottleneck that a static market estimate hides, such as limited integration staff or slow enterprise procurement.

6. Sensitivity, uncertainty and falsification

Market models become decision tools when they show which assumptions matter most. Holding N at 150,000, reachability at 40% and enablement at 50%, the following sensitivity table varies conversion and annual revenue per account. Every entry is an illustrative scenario in US\$ millions, not a confidence interval or forecast.

Paid conversion	US\$600 / year	US\$1,200 / year	US\$2,400 / year
10%	1.8	3.6	7.2
20%	3.6	7.2	14.4
30%	5.4	10.8	21.6

The table makes the commercial question explicit. Raising price may reduce conversion; the two inputs are unlikely to be independent. A high-price, high-conversion corner should not be labeled the expected case without evidence. Likewise, strong reachability through a partner may come with revenue sharing, restricted ownership of the customer relationship or concentration risk.

For an initial study, interval estimates should be empirical where possible. A conversion rate derived from a pilot needs its numerator, denominator, recruitment method and observation period. A convenience sample of enthusiastic partners should not be presented as representative of all practices. A confidence interval addresses sampling variability, not unmeasured selection bias.

Predeclare the thesis breakers

Three findings would materially weaken the proposed model: enablement costs exceed the customer's plausible benefit; the buyer with budget does not capture the benefit; or operational improvements do not persist after intensive pilot support ends. These are not peripheral execution issues. They can invalidate the commercial mechanism even when the broad social need remains large.

Research should therefore test willingness to pay with specific workflow and support commitments. A general expression of interest in innovation is weak evidence. A paid pilot, a signed budget allocation or a documented procurement process is stronger, although none guarantees renewal. Record objections and declined participation alongside successful conversations.

The decision rule should be written before the pilot: continue, narrow, redesign or stop based on net benefit, reliability, safety and observed demand. A research program that can only produce a positive conclusion is a marketing program. A transparent research process earns credibility by preserving negative results and making changes traceable to them.

7. Where interoperability creates value

The proposed opportunity is a portfolio of cross-organizational workflows. Each workflow has a sender, receiver, subject, permission boundary, expected disposition and economic owner. This structure is more actionable than a generic promise to connect every system. It also permits a common technical foundation without assuming that every sector has identical rules.

Candidate workflow	Observable value event	First research obstacle
Referral coordination	Appropriate referral accepted and attended	Capacity, patient preference and status visibility
Intake portability	Verified information reused with less repetition	Identity, provenance and recipient acceptance
Eligibility / authorization	Complete request with timely disposition	Payer-specific requirements and exceptions
Care discovery	Suitable available service selected	Directory freshness and ranking incentives
Payment reconciliation	Correct settlement matched to service	Refunds, disputes and responsibility allocation

An integration should be valued by the failure it removes. For example, a referral adapter that only transmits a message may save a few minutes but leave the receiving queue unchanged. A closed-loop workflow that detects rejection, identifies a permitted alternative and returns an accountable status has a different value proposition. The latter also imposes greater operational and governance obligations.

The buyer-beneficiary relationship deserves explicit study. A provider may perform the onboarding while a payer benefits from lower downstream utilization. A patient may experience reduced friction while a clinic bears support costs. If the product depends on one party subsidizing another, the contract must make that exchange sustainable. Social value does not automatically become a viable revenue stream.

A proposed portfolio discipline

Prioritize workflows using evidence quality, frequency, avoidable burden, feasible permissions, implementation cost and the severity of failure. Do not collapse these dimensions into an unexplained score. A promising high-volume workflow can be rejected because its errors are difficult to detect or because a necessary data-sharing relationship is unavailable.

An accountable coordination layer spanning clinical and adjacent care contexts is one proposed mechanism for expanding access. Its contribution remains an empirical question. It should be evaluated against simpler alternatives, including improved directories, staff training, existing vendor configuration and basic electronic forms. Innovation is strongest when it creates a new feasible workflow, while still demonstrating that the added complexity is justified.

8. International, regional and specialty lenses

The proposed research taxonomy uses three independent dimensions. **International** compares systems and cross-border constraints. **Regional** examines implementation within a country, state or local service area. **Specialty** examines workflows with distinct clinical or operational requirements. A report can carry all three tags rather than being forced into one exclusive folder.

Lens	Example research question	Required local evidence
International	Which institutions control exchange and reuse?	National governance, identity and procurement structures
Regional	Where do referrals fail in a defined service area?	Local capacity, connectivity, languages and transport
Specialty	Which data determine an acceptable handoff?	Specialty-specific content, urgency and review rules

International expansion should be staged around a reproducible operating model rather than a list of countries. A region with strong digital infrastructure may still require substantial localization of contracting, terminology and reimbursement. A less digitized region may have a more immediate need for assisted navigation but a weaker direct software purchasing channel. Neither condition can be inferred from national income alone.

Specialty analysis should focus on task differences. Behavioral health may place distinctive constraints on sensitive records. Rehabilitation may depend on repeated appointments and referral continuity. Home-based services introduce travel and workforce availability. Dental and personal wellness may involve different payment patterns. These are research hypotheses to investigate locally, not universal descriptions of every organization in those categories.

The website research library can reflect this structure through a topic title, geography tag, specialty tag, evidence cutoff, document status and a plain-language abstract. Readers should be able to distinguish an empirical study, a narrative review, a technical proposal and an illustrative economic model before opening the PDF. Updating a category should not silently change an older report's findings.

Local validation before transfer

A successful pilot is evidence about its participants and context. Before transferring it, compare infrastructure, staffing, legal permissions, language, service capacity and buyer incentives. Where a determinant differs, identify which part of the intervention must change and which outcome needs remeasurement. This makes expansion a sequence of testable adaptations rather than an unsupported claim of universal applicability.

9. Research design for a credible first market

The proposed first study has three linked components: an account census for one bounded segment, a workflow observation study and a prospective commercial pilot. The purpose is to produce a defensible local estimate before extrapolating. Each component should use the same account definitions so that evidence can move from the registry to the operating model without changing units.

The account census starts with authoritative business-statistics definitions. US County Business Patterns describes employer establishments; Nonemployer Statistics covers a separate business population. These sources support a counting method but have not been extracted into a new market count in this paper.^{[4][5]} A future extraction should preserve release year, industry code version, geography, exclusions and suppression handling.

Stage	Proposed evidence	Decision output
Registry construction	Deduplicated organizations and establishments	Bounded eligible-account denominator
Workflow observation	Time, failures, handoffs and responsible roles	Specific intervention and comparator
Feasibility screening	Contracts, data access and integration effort	Serviceable subset and enablement cost
Prospective pilot	Outcomes, burden, safety and paid demand	Continue / narrow / redesign / stop

Observation should include both completed and failed cases. Sampling only successful referrals or paid transactions creates survivorship bias. A lightweight event log can record initiation, missing information, acceptance, rescheduling, abandonment and closure without collecting unnecessary clinical detail. Researchers should reconcile event logs with staff observations to detect work occurring outside the digital system.

The pilot should compare against an explicit counterfactual. Depending on volume and contamination, that could be randomized rollout, matched sites or an interrupted time series with sufficient pre-intervention observations. The choice must be justified before outcomes are inspected. A before-and-after chart without control for staffing or seasonal change cannot establish causation by itself.

Commercial evaluation should continue after intensive implementation support declines. Record renewal, support demand, integration maintenance and actual resource redeployment. Freed staff time is not automatically a cash saving: it may improve service quality, reduce overtime or create capacity. Those outcomes are valuable but should be reported in their own units.

10. A public-interest research program

Researchers and implementing organizations could organize the opportunity as a sequence of evidence gates. The first gate verifies a need and buyer; the second verifies safe execution; the third verifies sustained value; the fourth verifies transferability. This sequence supports ambitious innovation without requiring every technical and commercial uncertainty to be solved simultaneously.

Gate	Required artifact	Release criterion
Need and denominator	Account register plus workflow baseline	Defined buyer, burden and alternative
Feasibility	Data-flow map, contracts and integration test	Permitted, observable and recoverable execution
Value	Prospective outcome and cost analysis	Benefit exceeds total operating burden
Transferability	Replication in a materially different setting	Performance survives changed conditions

The proposed analytics function should act as a research and monitoring service, not as an unrestricted autonomous brain. It can maintain evidence freshness, detect shifts in failure rates, compare cohorts and recommend investigations. Changes to clinical, financial or permission-sensitive workflows should pass through explicit review. The function's credibility depends on preserved data provenance and clear authority boundaries.

Research governance should include a named owner for each claim. Market claims need a statistical owner; product claims need an engineering owner; clinical interpretations need qualified review; legal applicability needs jurisdiction-specific review. A single impressive PDF should not be allowed to certify all four domains. This paper supplies a framework for those reviews rather than asserting that they have occurred.

Portfolio sequence

A plausible sequence is to start with read-only discovery and status visibility, then introduce supervised intake reuse, and only later test consequential writes or financial actions. That ordering is a proposal based on reversibility and observability. It is not a universal prescription: a region with mature partners may support a different entry point, provided the evidence gates remain explicit.

The most valuable eventual publication would report what failed as well as what improved. A transparent record of rejected workflows, unsuccessful onboarding and inequitable outcomes would help others judge the boundaries of the intervention. That discipline can establish whether coordination expands practical access, whose needs remain unmet and which institutional conditions are necessary for equitable implementation.

Appendix A. Reproducible market data specification

The following schema is proposed for the next empirical sizing exercise. It is intentionally designed to separate identity, eligibility and commercial assumptions. A completed dataset should include a versioned data dictionary, source snapshots or permitted archival references, and a transformation log. No real organization-level dataset is supplied by this working paper.

Field group	Minimum fields	Validation rule
Source provenance	Publisher, release, URL, extraction date	Every imported row traces to a source
Entity identity	Jurisdiction, identifier, legal name, parent	IDs unique within their issuing system
Operating location	Establishment ID, geography, service type	Separate location from purchasing entity
Eligibility	Inclusion flag, rule, exclusion reason	Mutually interpretable categories
Workflow readiness	Systems, access, staff owner, permission	Unknown distinct from absent
Commercial status	Contactability, buyer, price, decision date	Interest distinct from signed purchase
Quality	Duplicate confidence, missingness, reviewer	No silent imputation of key fields

Reproduction sequence

Freeze the source release and industry classification. Import organizations and establishments into separate tables. Standardize identifiers before names. Review candidate duplicates with a documented rule. Assign each eligible establishment to a buying account where evidence permits; retain unresolved relationships as unknown. Compute counts by geography and segment before calculating any global roll-up.

Publish both gross and deduplicated totals with the reconciliation. If a parent organization has ten locations, the report may show ten service locations and one buying account, but it should never add those values. Where a business operates across specialties, use a primary classification for mutually exclusive totals and secondary tags for analysis. Explain how multi-category records are handled.

Missingness should be reported as a finding. A missing endpoint does not prove no endpoint exists; an unverified organization is not automatically inactive. Perform sensitivity analysis under explicit assumptions about unresolved records. Do not hide uncertainty by assigning every unknown to the most commercially favorable category.

The analytical package should contain formulas, units, scenario inputs and an executable calculation file. A reviewer must be able to reconstruct the headline from the frozen inputs. Any later change in a source, definition or assumption should trigger a new model version, with the previous estimate retained for comparison.

Appendix B. Claim release & limitations

This paper is a targeted narrative synthesis, not a systematic review. Sources were selected for relevance and authority, including government statistics, standards documentation and the supplied planning file. No exhaustive search protocol, registered review protocol or formal publication-bias analysis was performed. The proposed market and evaluation models are original analytical constructions requiring empirical validation.

Proposed external claim	Evidence threshold	Current status
Healthcare is a large economic sector	Official expenditure statistics with year and geography	Supported as contextual statement
A specific global provider count is addressable	Compatible unit definitions and deduplicated data	Not established
An intervention can integrate a named system	Authorized access plus verified end-to-end test	Not established by this paper
A workflow saves money	Comparative full-cost analysis	Proposed study only
The product improves access or equity	Patient-level outcomes with subgroup analysis	Proposed study only
A scenario is company revenue guidance	Validated pipeline and operating assumptions	Not provided

The expenditure baselines have different observation years. The interoperability evidence is specific to US nonfederal acute care hospitals. The account-model inputs are hypothetical. Business-statistics documentation was used to define a future extraction, not to produce new official counts. These limitations are structural and should travel with any excerpted conclusion.

No original clinical or patient data were analyzed. Proposed architectures do not establish deployed interfaces, permissions, contracts, security controls or customer relationships. Publisher branding does not imply endorsement by cited institutions.

Recommended next empirical deliverable

Commission one bounded regional opportunity study with a frozen account denominator, observed workflow burden, verified enablement costs and actual buyer interviews. Publish the method and negative findings. Use that study to replace the hypothetical inputs in Section 5. Only then consider a broader estimate based on explicitly justified transfer assumptions.

Closing proposition. The strongest opportunity narrative connects a large social need to a small set of measurable, repeatable mechanisms. A care hub can be ambitious in scope and exacting in evidence at the same time. The work ahead is to demonstrate which connections produce better care journeys, who pays for them and whether the improvement persists at scale.

Appendix C. The sextillion scenario

The preliminary mathematical research document located for this review is the 19-page September 14, 2026 revision titled The mathematics of a universe of care. Its central figure is **1.56 sextillion theoretical combinations**, rather than 1.65 sextillion. The source explicitly treats the inputs as planning assumptions and the combinations as unconstrained possibilities.^[7]

The arithmetic is straightforward in short-scale English: 6.5 billion people × 800 provider-taxonomy labels × 300 million hypothetical products = 1.56×10^{21} combinations. One sextillion equals 10^{21} . The result is a Cartesian-product count of labeled configurations, not a count of customers, transactions, clinically valid pathways or dollars.

Dimension	Source scenario input	Interpretation
People	6.5 billion	Planning population, not verified platform reach
Provider labels	800	Scenario assumption; the retrieved NUCC 26.1 file has 883 unique codes ^[8]
Products	300 million	Hypothetical catalogue, not live inventory
Product of dimensions	1.56 sextillion	Unconstrained configuration count

The taxonomy count was independently checked against the NUCC version 26.1 CSV: 883 data rows and 883 unique code identifiers. Retaining 800 is permissible only as an explicitly hypothetical parameter. Replacing it mechanically with 883 gives 1.72185 sextillion raw combinations, which still does not validate the population or product assumptions.^[8]

From mathematical possibility to feasible action

Let X be the set of raw configurations and let $F(x)$ indicate that a configuration satisfies all required constraints. The feasible set is the subset of X for which $F(x) = 1$. Its size is the sum of those indicators, not automatically the product of separate eligibility percentages. Constraints can be correlated: a provider's specialty, geography, authorization and available products are related.

The economically relevant subset is smaller still: a feasible configuration must correspond to a real need, a chosen action, an available service and a sustainable payment arrangement. The same person may appear in many configurations without creating multiple buyers. Several syntactically different configurations may describe the same practical journey.

The source is useful as an illustration of why search, ranking and orchestration become complex. It is not a market-size estimate. Research can connect the two topics by studying how effectively a system identifies and executes useful feasible options within a large possibility space, while keeping commercial sizing grounded in observed accounts and transactions.

Democratization of care: an evaluation framework

Interoperability as shared access infrastructure

Interoperability matters to democratization when it reduces the practical costs of participating in care. A technical connection can reduce duplicate paperwork, help a receiving organization understand a referral or permit a person to move between providers without rebuilding their history. These mechanisms need outcome evidence: the existence of an API is an intermediate capability, not the social result.

Large organizations may have resources to purchase interfaces, negotiate contracts and maintain specialized staff. A system that requires the same fixed investment from a small community provider may deepen concentration. Evaluation should therefore measure the distribution of onboarding costs, maintenance burden and access terms across organizational sizes. Open specifications can help, but pricing, documentation, portability and operational support determine who can participate in practice.

Unmet need is not identical to paying demand. Public financing, cross-subsidy, community ownership and shared infrastructure may support valuable connections that cannot sustain a conventional per-user fee. A public-interest appraisal should compare these institutional options and publish who funds the service, who captures savings and whether the arrangement remains affordable after any pilot subsidy ends.

Dimension	Proposed measure	Interpretation safeguard
Participation	Share of eligible small or community providers with a usable connection	Count operational connections, not directory listings
Access	Completed appropriate care among all eligible people	Include offline and unsuccessful journeys in the denominator
Cost distribution	Patient burden and provider integration cost by group	Report who pays and who receives savings
Contestability	Cost and time to change service or export usable records	Test actual portability, including assistance

These measures are a proposed evaluation framework, not established findings about an existing service. Report baseline conditions, uncertainty, excluded populations and adverse results. A credible study can conclude that an intervention is useful, ineffective or inequitable; democratization is the question being tested, not a benefit assumed in advance.

References & evidence access

Numbered references identify source material supporting nearby factual statements. Interpretation, proposed architecture and labeled scenarios are original synthesis. Web sources were accessed for this review; the evidence cutoff is October 2, 2026. Report dates and underlying observation years are retained separately. No exhaustive database search or formal risk-of-bias meta-analysis was performed.

[1] World Health Organization. Global spending on health: Emerging from the pandemic. 2024; observations through 2022.

Official report. ISBN 978-92-4-010449-5. Global expenditure baseline; not a software-market estimate.

<https://www.who.int/publications/i/item/9789240104495>

[2] Centers for Medicare & Medicaid Services. National Health Expenditure Data: NHE Fact Sheet. 2024 historical expenditure series; retrieved for 2026 review.

Official statistical summary. Dollar values are nominal US dollars and the geography is the United States.

<https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/nhe-fact-sheet>

[3] ASTP/Office of the National Coordinator for Health Information Technology. Interoperable Exchange of Patient Health Information Among U.S. Hospitals: 2023. 2024 publication; 2023 survey observations.

Official survey-based descriptive brief. Universe: US nonfederal acute care hospitals; not all providers or countries.

<https://healthit.gov/data/data-briefs/interoperable-exchange-patient-health-information-among-us-hospitals-2023/>

[4] US Census Bureau. County Business Patterns: About this Program. Program documentation, accessed for 2026 review.

Primary statistical-method documentation. Employer establishments; this series did not extract a new establishment count.

<https://www.census.gov/programs-surveys/cbp/about.html>

[5] US Census Bureau. Nonemployer Statistics: About this Program. Program documentation, accessed for 2026 review.

Primary statistical-method documentation. Nonemployer business units require separate treatment from employer establishments.

<https://www.census.gov/programs-surveys/nonemployer-statistics/about.html>

[6] User-supplied Careverse planning document. Global Interoperability Opportunity by Provider Universe 1(3).pdf. Undated; supplied source reviewed October 2026.

Internal seven-page planning source. No supporting citations or reproducible count methodology identified.

Quoted ranges are unverified hypotheses, not external evidence.

[7] Careverse internal research revision. The mathematics of a universe of care. September 14, 2026; 19 pages.

Retrieved source: Careverse_Care_Flow_Math_Research_Revision(2).pdf. Pages 2 and 16-17 present the 1.56 sextillion scenario. Planning inputs and unconstrained combinations, not validated market demand.

[8] National Uniform Claim Committee. Health Care Provider Taxonomy Code Set, version 26.1. July 1, 2026; downloaded October 2, 2026.

Audit of the downloaded CSV: 883 data rows and 883 unique Code values. Count does not imply independent provider categories or a global care-sector classification.

https://www.nucc.org/images/stories/CSV/nucc_taxonomy_261.csv