

InXite Evidence Timed Outcome Journey Intelligence Master White Paper on the InXite Method and the Journey Cost Twin

A governed method for turning fragmented evidence into timed pathways actions outcomes and financial consequences

Controlled Counsel and Product Review Edition | 16 September 2026

Published by InXite.Health, a division of InXite Data Systems LLC.

Central thesis

InXite Evidence-Timed Outcome Journey Intelligence is the generalized proprietary method. It converts authorized evidence into a time-aware view of current state, plausible future pathways, actionable decision windows, expected outcomes and cost or value. In healthcare, the Journey Cost Twin applies this method to a person's health, coverage, preferences, care choices and financial exposure. ClearPrice closes the loop by carrying qualified price options through verification, scheduling or fulfillment, service, payment, reconciliation and measured savings.

Executive Summary

Health information is fragmented across electronic health records, health plans, devices, applications, consumer-entered information and, when a person explicitly chooses to share it, selected social content. Connecting these sources is necessary, but connection alone does not help a person decide what matters now, what may happen next, what choices are available or what those choices may cost.

InXite Health proposes the Journey Cost Twin as a novel application of its governed Whole-Person Digital Twin. The strengthened method maintains an Adaptive Journey Graph of possible future states and decision points, estimates how reasonable actions may affect those pathways, and attaches a personal Financial Twin to each meaningful event. LOLA by InXite provides the conversational explanation and orchestration layer. ClearPrice provides the price discovery, comparison and price-securing capability that converts a modeled opportunity into a verified consumer action.

The strengthened method separates current-state estimation, journey forecasting, action-impact modeling and personalized optimization. It distinguishes evidence uncertainty, model uncertainty and future uncertainty. It uses a Personal Utility Model to rank clinically reasonable choices according to the person's priorities. It calculates the value of missing information, detects divergence between expected and actual journeys, and learns from decisions, actions and measured outcomes.

The initial product can be demonstrated with synthetic data and bounded rules, scenario simulation and transparent cost logic. Production forecasting requires retrospective model development, clinical and actuarial validation, calibration, bias analysis, security and privacy controls, regulatory review and prospective measurement. InXite should describe early forecasts as modeled possibilities and cost opportunities, not diagnoses, guarantees or validated savings.

Core Recommendation

Decision	Recommended position
Product identity	Use Whole-Person Digital Twin for governed context and Journey Cost Twin for the combined Adaptive Journey Graph, Action Impact Engine and Financial Twin.
Consumer promise	Help each person understand what may be ahead, what could change it and what each choice may cost.
Commercial entry	Lead with ClearPrice, then use the connected record and LOLA to expand from immediate savings into preventive and proactive planning.

Decision	Recommended position
Technical approach	Combine rules, temporal models, Bayesian updating, causal evidence, scenario simulation, personal utility and claims-based cost models. Use large language models for extraction, explanation and conversation rather than as the sole prediction engine.
Claims discipline	Separate demonstrated capability, modeled output, validated performance and realized savings.

1 The Problem Is Rationalization

Most health platforms organize information around a single source or single moment. An EHR describes delivered care. A plan describes coverage and adjudicated claims. A wearable describes selected measurements. A health application describes behaviors within its own workflow. Consumer-entered information describes lived experience. Social posts may reveal interests, concerns or intent, but only when the person explicitly selects them for use and only with careful evidentiary labeling.

A unified feed can place these items beside one another without resolving identity, meaning, time, duplication, conflict, relevance or permission. That produces more information but may not produce better decisions. InXite's differentiation begins after connection, when the platform determines what each datum means, whether it is current, how it relates to other evidence and what remains uncertain.

Five Questions Every Important Output Must Answer

- What authorized sources were used
- What InXite resolved or reconciled
- What the system learned about the person
- What the system sees ahead and with what confidence
- What the person can do next and how that action may change health and cost

This structure should appear in the product, not remain hidden in technical documentation. It gives the consumer a practical reason to trust, question or correct the output. It also creates an auditable chain from source evidence through inference to recommendation and action.

Source	Useful contribution	Required treatment
EHRs and health information exchanges	Conditions, medications, allergies, observations, procedures, notes and care plans	Preserve origin, author, date, status and clinical context

Source	Useful contribution	Required treatment
Health plans and claims	Coverage, utilization, adjudication, network status and member cost history	Separate billed, allowed, paid and member responsibility amounts
Devices and health applications	Longitudinal activity, sleep, heart rate and other measurements	Retain device identity, measurement quality, cadence and gaps
Person directed information	Symptoms, goals, beliefs, preferences, barriers, readiness and outcomes	Attribute to the person and support confirmation or correction
Explicitly shared social content	Selected interests, concerns, community context or intent	Treat as reported or inferred context, never as clinical fact
Price and access data	Cash prices, negotiated rates, Medicare benchmarks, availability and discounts	Normalize service definitions, setting, locality, bundle contents and freshness

The Generalized InXite Method

Evidence-Timed Outcome Journey Intelligence is broader than a healthcare digital twin. It is a reusable computational and governance method for any environment in which evidence arrives from multiple sources, decisions occur within time windows, several outcomes remain possible and later results should improve the next recommendation.

Canonical Processing Loop

The method receives authorized evidence, preserves its origin and timing, rationalizes identity and meaning, estimates the best-supported current state, constructs plausible outcome journeys, identifies remaining decision windows, compares reasonable actions, models expected consequences, records the human decision and measures what actually happened. Selective recalculation updates only the outputs that depend on changed evidence, permissions, actions or outcomes.

Healthcare Embodiment

The Journey Cost Twin is the healthcare embodiment of the generalized method. It combines clinical, financial, behavioral, access and person-directed context while keeping evidence, inference, forecast and recommendation visibly distinct. The same generalized method can support enterprise, investor and other non-health outcome journeys without transferring healthcare-specific records or rules into those applications.

2 The InXite Journey Cost Twin

The Journey Cost Twin is a consumer-governed computational model that jointly forecasts personal health trajectories and their financial consequences. It does not attempt to predict one inevitable

future. It maintains a set of plausible pathways, updates their relative likelihood as evidence changes and explains which assumptions create the differences among them.




Figure 1 The Journey Cost Twin couples health and financial reasoning inside a continuous learning loop

The Novel Combination

Individual components of the design are established techniques. Interoperability standards can move health and financial data. Temporal and survival models can estimate time to an event. Bayesian methods can update beliefs as evidence changes. Scenario simulation can compare possible futures. Price files and claims can estimate cost. Conversational AI can explain complex information. The proposed innovation is the governed orchestration of these techniques around one person and one evolving decision process.

The method binds each forecasted event to six objects: evidence, probability or likelihood band, time horizon, cost shadow, intervention opportunities and uncertainty. A pathway is therefore more than a predicted diagnosis or cost total. It is an inspectable sequence of possible events and decision points that a person and care team can evaluate.

Journey object	Purpose
Current state	The best supported representation of present health, context, goals and constraints
Possible event	A bounded event that may occur within a defined time horizon
Pathway	An ordered set of possible events under stated assumptions
Cost shadow	The expected financial consequence of an event or pathway for the person and other stakeholders

Journey object	Purpose
Intervention	A choice or action that could change likelihood, timing, burden or cost
Learning signal	A later observation that confirms, weakens or changes the model


Adaptive Journey Graph

The Adaptive Journey Graph converts a set of scenarios into a living decision structure. Nodes represent current or possible future states. Edges represent temporal progression, clinical dependencies, consumer choices or modeled intervention effects. Each node and edge retains its evidence, assumptions, time horizon, confidence and financial consequences. The graph updates selectively when new evidence affects only part of the journey.

Graph element	Required contents
State node	Health status, functional status, goals, access conditions and supporting evidence
Possible event	Time horizon, likelihood band, evidence, uncertainty and downstream dependencies
Decision point	Reasonable options, professional boundary and information required
Action edge	Action, targeted driver, expected effect, effort, burden and monitoring plan
Financial edge	Consumer, sponsor and total cost distributions under stated assumptions
Learning signal	Observed result, divergence from forecast and resulting graph update

3 Evidence Aware Rationalization

Rationalization is the control layer that makes downstream prediction defensible. The system first establishes identity and authorization, then normalizes terminology, aligns time, removes duplicates, detects contradictions and classifies the evidentiary state of each material element. The original source is never erased by a consumer correction or model inference.



Every state carries source time permissions freshness confidence and correction history

Figure 2 Evidence states prevent reported observations and model outputs from silently becoming clinical facts

Rationalization Pipeline

Stage	Technique	Output
Authorization	OAuth, consent policies, purpose and recipient checks	Permitted source scope and duration
Identity	Deterministic and probabilistic matching with review thresholds	Linked records with unresolved identities exposed
Semantic normalization	FHIR profiles, code systems, units and service taxonomies	Comparable clinical and financial concepts
Temporal alignment	Event time, recorded time, effective period and freshness rules	Longitudinal sequence with stale data identified
Conflict resolution	Source hierarchy, recency, context and explicit reconciliation	Resolved element or visible conflict state
Evidence qualification	Known, person reported, device observed, inferred, conflicting or unknown	Inspectable evidence graph
Feature derivation	Transparent transformations with lineage	Model features that can be traced to source evidence

HL7 FHIR provides standard resources for clinical and financial exchange, including Provenance and ExplanationOfBenefit. InXite should retain these native semantics where available and add its own evidence, confidence, permission and correction metadata when source systems do not provide enough detail. Provenance is part of the model, not an administrative afterthought.

4 A Living Whole Person Model

The Whole-Person Digital Twin is the governed context layer beneath the Journey Cost Twin. It is not a perfect replica of a person. It is an evolving evidence graph that connects clinical state, behavior, goals, access, relationships, decisions, actions and outcomes.

Domain	Illustrative contents	Why it affects the journey
Clinical state	Conditions, medications, allergies, labs, imaging, procedures and notes	Defines professional evidence and care dependencies
Behavior and lifestyle	Activity, sleep, nutrition, substance use, routines and adherence	Changes exposure, feasibility and response
Goals and priorities	What the person wants to improve, avoid or preserve	Determines which tradeoffs matter
Beliefs and preferences	Treatment preferences, wellness orientation and communication needs	Shapes acceptable options without changing evidence
Access and economics	Coverage, deductible, affordability, transportation and availability	Determines whether a recommendation is practical
Relationships and responsibilities	Care team, caregivers, dependents and work demands	Affects coordination and implementation
Decisions and actions	Options considered, choices made, actions started, deferred or declined	Separates recommendation failure from implementation failure
Outcomes and learning	Clinical, functional, experiential and financial results	Updates the model and tests whether the plan worked

Personalization Without Evidence Distortion

Personal priorities should change ranking, timing, explanation and plan design. They should not change the underlying clinical evidence. For example, a person's desire to avoid surgery may cause the system to present conservative options first, while still showing the evidence, limitations and circumstances that could make surgery appropriate. The same distinction applies to cost, distance, time away from work, caregiving responsibilities, faith, wellness practices and longevity goals.

Consumers must be able to confirm, correct, defer or reject material assumptions. A correction should create a new attributed state rather than overwrite the source. This supports trust, auditability and better future model performance.

Personal Utility Model

The Personal Utility Model represents the tradeoffs that matter to the individual. Clinical evidence and safety rules first determine which options are reasonable. The utility model then ranks those options using the person's priorities, such as expected health benefit, symptom relief, independence,

longevity, cost, travel, speed, treatment burden, time away from work, caregiver impact and willingness to act. A change in priority may change the ranking, but it does not change the underlying evidence.

The product should make sensitivity visible. For example, it can explain that one option offers the strongest price and clinical fit, while another becomes the better personal fit when proximity or faster access receives more weight. The consumer remains able to adjust priorities or reject the proposed ranking.

5 Forecasting Techniques

A credible first release should use a hybrid forecasting stack. No single model is likely to handle heterogeneous health evidence, sparse individual histories, changing coverage, local price variation and the need for clear explanations. Each model should have a bounded intended use and an explicit fallback when data are insufficient.

Technique	Role in InXite	Primary control
Rules and pathway knowledge	Encode clinical dependencies, care sequences, exclusions and known decision points	Versioned content, professional review and jurisdictional scope
Temporal feature engineering	Convert longitudinal observations into recency, trend, persistence and change features	Lineage, missingness flags and leakage prevention
Time to event models	Estimate whether and when a bounded event may occur	Calibration by horizon and subgroup
Bayesian updating	Revise event likelihoods as new evidence arrives	Visible prior assumptions and sensitivity analysis
Gradient boosting or neural models	Detect nonlinear patterns when validated data volume supports them	External validation, explainability and drift monitoring
Scenario simulation	Compare current, proactive, delayed and alternative pathways	State assumptions and avoid causal claims without evidence
Causal methods	Estimate intervention effects when study design and data support the question	Confounding analysis and target trial discipline
Large language models	Extract concepts, reconcile narrative context and generate grounded explanations	Retrieval, citations, output constraints and human escalation

Health Pathway Scenarios

The consumer should see a small number of useful scenarios rather than an exhaustive tree. A typical view can include the likely current trajectory, a proactive or preventive pathway, delayed

action, higher risk when clinically relevant, an alternative care pathway and a wellness or longevity pathway. Scenarios should appear only when the evidence and intended use support them.

Each modeled event should show its time horizon, likelihood band, supporting evidence, missing information, confidence, factors that could change the forecast and the next decision or intervention opportunity. Exact probabilities should be shown only when the model is sufficiently calibrated and the number improves the decision.

Four Intelligence Functions

Function	Question	Required discipline
Current state estimation	What is most likely true now	Resolve source, time, conflict and missingness
Journey forecasting	What may happen next	Calibrate by population and time horizon
Action impact modeling	What could change the journey	Separate causal evidence from association
Personalized optimization	Which reasonable option best fits this person	Apply safety constraints before personal preferences

Action Impact Engine

The Action Impact Engine links each recommendation to the journey driver it is intended to influence. It records the proposed action, supporting evidence, expected health and financial effects, time to expected benefit, consumer effort, burden, confidence, monitoring plan and conditions requiring professional review. When causal evidence is insufficient, the system describes an action as a hypothesis or scenario rather than claiming that it will produce the modeled outcome.

Three Forms of Uncertainty

Uncertainty	Meaning	Consumer expression
Evidence uncertainty	Information is missing, old, conflicting or measured inconsistently	What InXite needs to confirm
Model uncertainty	The model or validation population may not fully represent this person or setting	How reliable the estimate is
Future uncertainty	Later choices, clinical developments, coverage and prices may change	What could make the journey different

6 Joint Health and Cost Forecasting

The cost model should operate at the same event level as the health model. When a pathway contains a possible test, visit, procedure, medication, supplement, device or service, InXite attaches a cost shadow that represents the best available personalized financial estimate under stated assumptions.

At its simplest, modeled expected cost is the sum of each event likelihood multiplied by that event's personalized cost estimate. Production models should retain distributions and dependencies rather than rely on one point estimate. A hospitalization and its follow-up services, for example, are correlated events and should not be treated as independent.

Expected pathway cost = sum of event likelihood x personalized event cost

Financial Twin Components

Component	Interpretation
Market benchmark	Local median price for a matched service definition and setting
Medicare benchmark	Locality and setting appropriate Medicare amount when the service is covered and comparable
Quoted or secured price	A provider, vendor or pathway quote with date, scope and conditions
Plan economics	Negotiated or allowed amount, deductible status, coinsurance, copay and exclusions when available
Consumer responsibility	Estimated out of pocket amount under the selected payment route
Sponsor or payer cost	Expected plan or program liability when relevant and authorized
Ancillary burden	Travel, time, caregiving, missed work or other costs the person chooses to model
Uncertainty range	Variation caused by incomplete coding, bundle contents, complications, coverage or price freshness
Administrative pathway	Eligibility, prior authorization, network and payment-route conditions that may affect access or cost

The forecast should distinguish price from total cost of care. A lower procedure quote can still create higher total cost if it excludes anesthesia, facility, pathology, medication, travel or follow-up. ClearPrice should therefore normalize bundle contents and show what is included, excluded and still unknown.

The Financial Twin extends the event-level cost shadow across the likely care sequence. It can compare cash, insurance, HealthShare and other authorized payment routes; represent deductible and accumulator status; model authorization and network conditions; and estimate consumer, sponsor and total cost distributions. Forecasted administrative outcomes must remain separate from verified eligibility, authorization and adjudication.

7 ClearPrice as the Closed Loop Action and Financial Evidence Layer

ClearPrice is InXite's primary consumer entry point and largest near-term revenue opportunity. It solves an immediate problem, demonstrates measurable value and creates a natural reason for the person to connect coverage, records and preferences. It also supplies the verified price and access data needed to turn a forecast into action.

Locked Comparison Standard

Every qualified option should be evaluated against the same procedure-level benchmarks. The screen keeps the local median and locality-adjusted Medicare benchmark at the top, then repeats the four core measures for every option. Quoted price receives the strongest visual emphasis. Savings and discounts use lighter shades. Green, yellow and red indicate objective price bands rather than provider quality or clinical appropriateness.

Procedure benchmark	Option level measures
Local median price	Quoted price
Medicare price	Discount from median
	Dollar savings from median
	Quoted price as a percentage of Medicare

Percentage of Medicare is calculated as quoted price divided by the matched Medicare benchmark multiplied by 100. The comparison must use a compatible code, locality, site of service, professional or facility component and bundle definition. Medicare amounts are reference points, not universal definitions of fair price, coverage or consumer liability.

From Opportunity to Realized Savings

State	What InXite may report
Modeled opportunity	Potential difference between a scenario or benchmark and an available option
Quoted savings	Difference supported by a dated quote with defined scope

State	What InXite may report
Secured price	Price reserved or contracted for the person under stated conditions
Realized savings	Difference supported by final transaction, claim or payment evidence

This claims ladder protects consumer trust and supports credible sponsor reporting. InXite should not report modeled opportunity as realized savings.

Price to Payment Journey

ClearPrice must persist one consumer-specific transaction from the qualified service definition through final financial reconciliation. The experience begins with a comparable bundle and dated option, then records verification, price security, appointment or fulfillment, prerequisites, service completion, claim or payment activity, variance and final consumer responsibility. The resulting evidence distinguishes a displayed estimate, quoted price, verified quote, secured price, completed service, final cost and realized savings.

Fulfillment Paths

The operating model supports distinct paths for insurance, self-pay, HealthShare or member sharing, and approved marketplace or product transactions. Each path applies its own eligibility, authorization, scheduling, payment and documentation requirements while retaining a common consumer view. Route-specific facts do not become universal facts, and a change of route produces a new version rather than silently rewriting the prior comparison.

Journey Cost Twin Update

A completed ClearPrice transaction updates only the affected financial and journey objects. A final price can recalibrate a local cost estimate, a delayed appointment can alter a decision window, and an incomplete prerequisite can change the expected sequence. The update retains the original forecast and selected action so InXite can explain the divergence instead of presenting the revised state as if it had always been known.

8 Next Best Learning

One of the most distinctive features of the proposed system is its ability to identify which missing fact, source connection or simple measurement would most improve the next decision. This is a value-of-information problem. The system estimates how much uncertainty would decrease, or how much the recommended action might change, if a specific missing item became known.

LOLA can then ask the smallest useful question rather than sending the consumer through a long generic assessment. Examples include confirming whether a medication is still being taken,

connecting current plan benefits, entering a home blood-pressure reading, clarifying the intended procedure location or identifying the person’s highest priority among price, distance, timing and preferred approach.

Candidate request	Possible information gain	Consumer burden
Confirm current medication use	High when the medication changes risk or pathway eligibility	One tap
Connect current plan	High when out of pocket cost and network ranking are uncertain	Authorization flow
Collect seven days of home readings	High when one clinical value may not represent the current state	Repeated measurement
Ask about affordability threshold	High when several clinically reasonable choices have different costs	One short answer
Request another record source	Variable depending on overlap, recency and missing data	Connection and consent

Decision Rule

The system should recommend a request only when expected decision improvement exceeds consumer burden, privacy cost and operational cost. Sensitive social content should never be requested by default merely because it might improve prediction. The person must choose what to share, for what purpose and for how long.

Value of information also improves product economics. It directs integration effort toward sources that are likely to change decisions and avoids collecting data merely because it is available.

Consumer Governed Experimentation

For appropriate low-risk wellness and behavior goals, InXite can help the person conduct a structured personal experiment. The system establishes a baseline, defines a bounded action and measurement period, monitors the selected outcomes and helps the person decide whether to continue, modify or stop. Examples may include sleep routines, activity, nutrition, stress management and adherence support. An individual result is personal learning and should not be presented as population-level clinical evidence.

Journey Divergence Detection

The system should compare the forecast with the observed journey. A missed action, unexpected device trend, new medication, delayed service, coverage change, different final price or faster-than-expected improvement can alter the pathway. Divergence detection identifies which assumptions no

longer hold, updates only the affected parts of the graph and explains the resulting change to the consumer and authorized care team.

9 LOLA and the Consumer Experience

LOLA is the persistent text and voice companion that translates the evidence graph, forecasts and price options into understandable decisions. LOLA should never conceal the underlying basis for an important recommendation. The consumer should be able to move from a plain-language answer to the source, the system’s reasoning, the uncertainty and the available action.

Consumer question	LOLA response structure
What changed	New evidence, the prior state, the current state and any unresolved conflict
Why are you showing this	Sources used, person-specific relevance, confidence and missing information
What may happen next	A small set of plausible pathways with time horizon and uncertainty
What can I do	Action choices ranked by evidence, fit, access, cost and the person’s priorities
What will it cost	Range, assumptions, plan route, cash route, median benchmark, Medicare benchmark and quotes
Can I change the plan	The effect of a new priority, constraint, decision or completed action

Signature Experience

The product experience should consistently show the sequence Connect, Rationalize, Personalize, Look Ahead, Act and Learn. The Home screen remains simple and price-led: LOLA and ClearPrice are followed by circular rows for My Health, My Care and My Coverage and Costs, then LOLA Insights and the fixed navigation. The sophistication belongs behind these familiar entry points and should surface only when it helps the consumer decide.

Across those entry points, the consumer should repeatedly see five questions answered: what InXite knows, what changed, what may be ahead, what could change the journey and which action currently offers the best personal value. My Journey shows the branches. My Choices exposes tradeoffs and sensitivity. ClearPrice converts a financial opportunity into a qualified option and secured price. LOLA explains and coordinates the next step.

Authority and Professional Boundaries

The recommended product language is clear: InXite recommends. You and your care team decide. High-consequence clinical questions should support review with a licensed professional. The product should identify urgent or emergency patterns through separately validated workflows and direct the person to appropriate care rather than continue an ordinary forecasting conversation.

Purpose Scoped Shared State

The consumer, provider, sponsor and operations teams should not receive identical views. A shared versioned state can support each role while policy services enforce purpose, recipient, duration, minimum necessary data and revocation. Consumer-directed context and derived forecasts remain unavailable for underwriting, adverse eligibility, individualized rate setting or risk selection.

10 Technical Architecture

The platform should separate source ingestion, governed evidence, model execution, conversational explanation and action services. This makes permissions enforceable, allows models to change without rewriting the source record and prevents a language model from becoming the system of record.

Layer	Core services and standards
Connection and exchange	FHIR R4 and SMART on FHIR for patient-authorized clinical access; payer APIs and CARIN profiles for claims and coverage; device and app APIs; controlled document ingestion
Canonical evidence	Longitudinal graph or hybrid store for normalized entities, events, provenance, consent, conflicts, versions and evidence states
Feature and model services	Versioned features, current-state estimator, Adaptive Journey Graph, Action Impact Engine, Personal Utility Model, Financial Twin, model registry and calibration service
Knowledge services	Terminology, procedure bundles, clinical pathway content, benefits rules, provider directory, market prices and Medicare fee schedules
LOLA orchestration	Retrieval grounded generation, structured tool calls, policy checks, explanation templates and escalation
Action services	ClearPrice quoting and price securing, scheduling, reminders, care coordination and plan tracking
Governance and operations	Consent, access policy, audit, security, observability, model monitoring, incident response and data retention

OpenAI and Distribution Strategy

The product should use multiple paths without making InXite dependent on one distribution channel. InXite should own the record, evidence and intelligence layer. OpenAI APIs can support LOLA extraction, explanation and conversation under InXite controls. A ChatGPT integration can provide consumer distribution and connection convenience. InXite can also pursue a direct OpenAI Health partnership. The core Journey Cost Twin, ClearPrice logic, governance, product experience and learning system remain InXite assets.

Model Execution Contract

- Every model call has a declared intended use, input schema, version and allowed consumer output.
- Every feature can be traced to authorized evidence and transformation logic.
- Every forecast carries horizon, confidence, calibration cohort and missingness indicators.
- Every high-impact recommendation uses a policy layer that can require professional review.
- Every action and later outcome can be linked back to the forecast and decision that preceded it.

Journey Model Card

Each predictive capability should maintain a professional-facing and consumer-facing Journey Model Card. It records intended and prohibited uses, target and validation populations, data sources, training and evaluation methods, performance by horizon and relevant subgroup, known limitations, update history, monitoring status and conditions requiring human review. The consumer version translates those attributes into the information needed to judge whether a recommendation is appropriate for the current decision.

11 Trust Privacy and Regulatory Design

The regulatory classification of any feature depends on its intended use, users, claims and behavior. Product, clinical, legal and regulatory leaders should review each use case before release. FDA's January 2026 clinical decision support guidance clarifies that many patient-facing or caregiver-facing software functions may remain within device oversight depending on their function. InXite should not assume that showing rationale alone resolves classification.

Risk area	Required design response
Clinical overreach	Bound intended use, avoid deterministic diagnosis and guarantees, and provide professional escalation
Hidden inference	Label inferred information, show rationale and allow correction
Bias and unequal performance	Measure calibration, error and benefit across relevant populations and access contexts
Privacy aggregation risk	Use purpose-specific authorization, minimization, revocation, retention limits and audit

Risk area	Required design response
Health app breach obligations	Maintain incident detection and notification processes appropriate to applicable law and role
Price mismatch	Validate code, locality, setting, bundle contents, coverage and freshness before comparison
Automation bias	Show alternatives, confidence and what would change the recommendation
Model drift	Monitor input, calibration, outcome and workflow drift with defined rollback criteria

NIST's AI Risk Management Framework organizes risk work through Govern, Map, Measure and Manage. InXite can translate those functions into model inventory, intended-use review, data and subgroup analysis, testing, incident response and ongoing monitoring. The framework is voluntary and is being revised; it should be used as a structured reference rather than represented as certification.

Health IT guidance for predictive decision support emphasizes source attributes, feedback capabilities and intervention risk management. FDA transparency principles likewise focus on communicating information that can affect risk and outcomes in ways suited to the intended user and workflow. InXite should treat the Journey Model Card, explanation interface and feedback loop as operational controls rather than publication-only documentation.

Health information may fall outside HIPAA when it is held by a consumer application that is not acting for a covered entity or business associate. Applicable privacy obligations can still arise under the FTC Health Breach Notification Rule, state law, contracts and consumer protection law. InXite should apply a strong common control baseline and then map obligations by product role and jurisdiction.

12 Validation and Readiness

The Journey Cost Twin can be prototyped now, but a compelling demonstration and a production forecast are different deliverables. Readiness should be reported by use case and claim, not as one platform-wide percentage.

Stage	Permitted demonstration	Evidence required to advance
Concept	Architecture, workflows and synthetic examples	Stakeholder review and intended-use definition

Stage	Permitted demonstration	Evidence required to advance
Functional alpha	Rationalized synthetic Twin, three scenarios, cost ranges, value-of-information prompt and ClearPrice action	Unit tests, traceability, security design and expert review
Retrospective prototype	Forecasts on de-identified historical data	Cohort definition, leakage controls, calibration, subgroup performance and error analysis
Prospective pilot	Bounded decision support with monitoring and human oversight	Protocol, governance approvals, workflow safety, user comprehension and incident process
Production use	Validated claims for named populations, settings and horizons	External or transport validation, regulatory determination, monitoring and change control
Outcomes expansion	Measured health, experience and financial impact	Prospective comparison, attribution method and realized-savings evidence

Minimum Model Evaluation Set

- Discrimination where appropriate, plus calibration by forecast horizon
- False-positive and false-negative consequences for the intended workflow
- Performance across age, sex, race, ethnicity, geography, disability, language, coverage and relevant clinical groups when data and law permit
- Stability under missing data, source delay, device gaps and contradictory evidence
- Cost error by service category, locality, site of service, payment route and bundle definition
- Consumer comprehension, correction behavior, action completion and professional escalation
- Realized health, experience and financial outcomes with attribution limitations stated
- Divergence-detection sensitivity, false alerts, time to detection and effect of graph updates
- Personal Utility Model stability, sensitivity and evidence that users understand why rankings change

13 Illustrative Consumer Journey

The following example is synthetic and demonstrates system behavior rather than clinical performance. A consumer searches ClearPrice for a knee MRI. With permission, InXite connects current coverage, a recent orthopedic note, prior imaging, device activity trends and the person's stated priorities. The system identifies inconsistent procedure descriptions and an unknown about whether contrast is required. LOLA asks one clarifying question because that answer materially changes comparable prices and the likely care sequence.

The Adaptive Journey Graph then displays three pathways: the likely current trajectory, a proactive conservative pathway and delayed action. Each pathway includes expected milestones, action-impact assumptions, the three forms of uncertainty and a Financial Twin cost range. ClearPrice

returns qualified options with the quoted price, discount from the local median, dollar savings and percentage of the matched Medicare benchmark. The consumer secures a price, schedules the service and chooses to share a focused Health View with the ordering clinician.




Figure 3 Synthetic scenario display values are illustrative and are not clinical or financial predictions

What the consumer sees	What the system records
Why the MRI comparison changed	New procedure detail, reconciled code and affected assumptions
Three possible journeys	Scenario version, event likelihoods, cost distributions and confidence
Best value options	Benchmark date, quote scope, locality, setting, availability and ranking criteria
Secure this price	Quote state, consumer authorization, transaction or referral state
Updated plan	Decision, action, due date, responsible participant and monitoring signal
Later outcome	Completion, final cost, consumer experience, health result and model update
Journey changed	Detected divergence, affected assumptions, revised pathways and explanation

14 Business and Intellectual Property Strategy

InXite’s commercial advantage should come from the system and its learning loop rather than a broad claim that the company uses a digital twin or artificial intelligence. The defensible asset is the governed method that links source evidence, an Adaptive Journey Graph, action-impact reasoning, personal utility, a Financial Twin, Next Best Learning, ClearPrice actions, divergence detection and later outcomes.

Asset	Potential defensibility
Evidence graph and ontology	Proprietary representation of source state, conflict, uncertainty, personal priorities, decisions, actions and outcomes
Journey event schema	Reusable linkage among evidence, probability, time, intervention and cost shadow
Adaptive Journey Graph	Structured representation of possible states, dependencies, decisions, action effects and learning signals
Action Impact Engine	Linkage of reasonable actions to modifiable drivers, expected effects, evidence and monitoring
Personal Utility Model	Transparent ranking of reasonable options according to consumer priorities and constraints
Financial Twin	Event and pathway cost distributions across payment routes, administrative conditions and stakeholders
Next Best Learning engine	Selection of the smallest useful question, source connection or measurement likely to improve a decision
ClearPrice normalization	Service matching, bundle comparison, benchmark logic, verification and price-securement workflows
Learning loop	Outcome-linked data showing which recommendations were accepted, implemented and effective
Consumer interaction design	Explainable progression from connection through rationalization, prediction, action and learning
Distribution strategy	InXite application, payer and provider channels, ChatGPT integration and potential OpenAI Health partnership

Potential patent protection requires a counsel-led prior-art and patentability analysis before public disclosure. Trade-secret protection may be more appropriate for model features, normalization methods, ranking weights, pricing operations and learning data. Contract, data-rights and brand strategy should support both approaches.

Commercial Flywheel

ClearPrice creates first engagement and immediate economic value. A successful comparison encourages record and plan connection. Better context improves personalization and journey forecasting. Forecasts create preventive and proactive actions. Completed actions and verified costs improve the learning system. Sponsors gain governed aggregate insight without receiving unrestricted access to individual records.

Revenue models can include price-securement or navigation fees, sponsor subscriptions, outcome-based arrangements, provider access services and premium consumer services. Each model needs clear conflict-of-interest controls so ranking remains faithful to the consumer's selected criteria and disclosed commercial relationships.

15 Implementation Priorities

Priority	Near term deliverable	Exit criterion
Define intended uses	Use-case cards for ClearPrice, journey forecasts, cost forecasts and LOLA recommendations	Clinical, legal, regulatory, product and security approval
Build canonical evidence	Source mapping, evidence states, provenance, time and permission model	Traceable synthetic end-to-end journey
Implement ClearPrice standard	Median, Medicare, quote, savings, discount and percentage-of-Medicare for unlimited qualified options	Benchmark and bundle test suite passes
Create Adaptive Journey Graph	Three transparent pathways across 30, 90 and 365 day horizons	Every state and edge has evidence, uncertainty and action logic
Build Action Impact Engine	Drivers, interventions, expected effects, burdens and monitoring plans	Causal and associational statements remain distinct
Add Financial Twin	Event and pathway cost ranges across payment routes and stakeholders	Reconciles against synthetic claims and quote examples
Add Next Best Learning	Ranked questions, connections and measurements with burden and privacy cost	System suppresses low-value or intrusive requests
Instrument divergence	Expected and actual state comparison, explanation and selective graph update	Closed-loop demonstration updates the Twin
Prepare validation	Data plan, model cards, evaluation protocol and governance process	Approved retrospective study design

Recommended Alpha Demonstration

The first alpha should tell one complete story rather than display disconnected features. It should begin with ClearPrice, rationalize at least five source types, show one meaningful conflict and one unknown, generate three branches in the Adaptive Journey Graph, expose an Action Impact assumption, attach Financial Twin ranges, ask one Next Best Learning question, let the consumer change a utility priority, secure a qualified price and show how divergence updates the journey after the action.

The alpha can use synthetic information and deterministic scenario rules. It should explicitly label those limitations. The purpose is to prove the interaction model, data contracts, governance and commercial flow before claiming predictive performance.

Conclusion

InXite has a credible path to a differentiated personal health intelligence platform. Its advantage is the ability to rationalize authorized evidence from clinical, financial, behavioral and person-directed sources, preserve what is known and uncertain, model several possible futures and connect each future to practical choices and costs.

The Journey Cost Twin makes that distinction visible. It combines the governed Whole-Person Digital Twin, Adaptive Journey Graph, Action Impact Engine, Personal Utility Model, Financial Twin, Next Best Learning, divergence detection, LOLA explanations and ClearPrice actions in one learning system. The person remains in control, and licensed professionals retain responsibility for clinical decisions.

InXite should proceed on multiple technical and distribution paths while retaining ownership of the record, evidence, journey, cost and outcome layers. The first goal is a transparent synthetic alpha. The next goal is disciplined retrospective and prospective validation. Market claims should advance only when the evidence advances.

Proposed Protectable Core

A governed system that converts authorized, time-qualified evidence into versioned current-state and pathway objects; identifies decision windows and missing information that may change an outcome; compares actions using expected outcome, personal utility and financial consequence; records the selected action; detects divergence between expected and observed events; and selectively recalculates dependent journey and cost objects. The healthcare embodiment adds a consumer-governed Journey Cost Twin and a ClearPrice transaction that carries a qualified option through route selection, fulfillment, service, payment, reconciliation and measured savings.

References

- 1 HL7 International FHIR Release 4 [View source](#)

- 2 HL7 International FHIR R4 Provenance [View source](#)
- 3 HL7 International FHIR R4 ExplanationOfBenefit [View source](#)
- 4 HL7 International International Patient Access Implementation Guide [View source](#)
- 5 HL7 International Bulk Data Access Implementation Guide [View source](#)
- 6 Centers for Medicare and Medicaid Services Interoperability and Prior Authorization Final Rule CMS 0057 F [View source](#)
- 7 Centers for Medicare and Medicaid Services Physician Fee Schedule Look Up Tool Overview [View source](#)
- 8 Centers for Medicare and Medicaid Services Hospital Price Transparency [View source](#)
- 9 Centers for Medicare and Medicaid Services Use of Pricing Information Published Under Transparency in Coverage [View source](#)
- 10 National Institute of Standards and Technology Artificial Intelligence Risk Management Framework 1.0 [View source](#)
- 11 National Institute of Standards and Technology Privacy Framework [View source](#)
- 12 US Food and Drug Administration Clinical Decision Support Software Guidance January 2026 [View source](#)
- 13 Federal Trade Commission Complying with the Health Breach Notification Rule [View source](#)
- 14 US Department of Health and Human Services Protecting the Privacy and Security of Health Information When Using Consumer Health Apps [View source](#)
- 15 Assistant Secretary for Technology Policy Health IT Responsible Artificial Intelligence in Health Care and HTI 1 [View source](#)
- 16 Assistant Secretary for Technology Policy Decision Support Interventions Certification Criterion Resource Guide [View source](#)
- 17 US Food and Drug Administration Transparency for Machine Learning Enabled Medical Devices Guiding Principles [View source](#)

About InXite Health

InXite Health is a consumer-first personal health intelligence and journey-orchestration platform. LOLA by InXite serves as the intelligent health companion and navigator. ClearPrice helps consumers find, compare and secure qualified prices for medical procedures, services, medications, supplements, alternative pathways and related products. InXite connects consumer, provider and population perspectives through governed evidence, coordination and learning.

Publication and Claims Notice

This controlled white paper presents a proposed product and technical framework for counsel, product and partner review. It is not approved for public technical distribution until patent counsel confirms the filing and disclosure sequence. It is not medical advice and does not assert that every

described capability, integration, model, outcome, price, discount, savings amount or economic effect is currently implemented, validated or commercially available. Synthetic examples are illustrative.

Published by InXite.Health, a division of InXite Data Systems LLC.