



ALTRAVIA

venture studio



STRADE ANTICHE, DESTINAZIONI NUOVE
ANCIENT ROADS, NEW DESTINATIONS
ESTRADAS ANTIGAS, NOVOS DESTINOS

ALTRAVIA

Trajectory, Not Status

An duine, ní an marc.

Human flourishing is not a state to be administered.
It is a trajectory to be navigated over time.

$$(2,2)_1 = (2,2)_2 \text{ but } T_1 \neq T_2$$



A working paper for the

Futura Innovation Conference
Limerick, 19–21 October 2026

and an invitation to the forty-five-minute
panel it exists to serve.



September 2026

Patrick Van Pelt
Founder

Sofia Menesello
Studio Associate

Francesca Pagani
Creative Director

altravia.studio

The Claim

Four people are about to receive the same answer.

Four people of the same age walk into a clinic on the same morning. Each has a blood-sugar reading of 5.7 percent, the number a laboratory calls the floor of prediabetes. The instrument prints the same word for all four.

- *One of them is stable and has been for six years.*
- *One is deteriorating and has been for eighteen months.*
- *One is recovering, and this is the best reading she has had since 2022.*
- *One has been measured exactly once. Nobody knows where he is going, including him.*

The number is correct in all four cases. The answer is useless in at least three of them.

I have spent thirty-five years inside instruments like that one, employee benefits, retirement plan committees, insurance, financial advisory, since 1991. The work is largely honest and the people in it are largely serious. And the measures we hand one another describe where someone sits today and say almost nothing about where they are going. We renew on a plan year. We stratify into tiers. We report a balance.

Then we act as though the person has been located.

They have not been located. They have been photographed.

This paper proposes that the unit of analysis is incomplete, that the improvements are available, and that the improvements have a price: an instrument built this way has to be willing to go quiet. The forty-five minutes in Limerick exist to consider that price.

WHAT THIS DOCUMENT CLAIMS, AND WHAT IT DOES NOT

It claims that trajectory is the right unit of analysis, and that an instrument should say plainly when it cannot supply one.

It claims no built instrument, no validated measurement, and no evidence that any intervention changes a trajectory.

The four people above are constructed cases, not patients. They can be shown without anyone's clinical permission, and they are evidence about a standard rather than about anybody's health.

Where a claim is untested, this document says so in the same sentence. The Status table lists every material claim and its status.

The Unit of Analysis

Same position, different trajectories.

Two cars cross the same mile marker at the same instant. One is nearly out of fuel and slowing. One was just filled and is speeding up. The marker records the equality. It cannot record the difference.

Two passengers sit in the same row of the same aircraft. One is flying to bury a parent. One is flying to meet a newborn grandchild. Same position. Not the same journey.

Institutions evaluate what is easy to count: activities completed, rules followed, money spent, conditions recorded at a point in time. Whether people are actually moving toward outcomes worth reaching is harder to count, so it tends not to be counted. The substitute gets chosen because it is countable — and then the substitute quietly becomes the thing. The trouble starts when we mistake the measurement for the person.

Two losses, worth naming separately.

Aggregation loss. Position is already a composite - health, money, capability, resilience, whatever the institution chose to include - compressed into a present state. Different underlying components produce the same composite number.

Kinematic loss. That composite is then observed without its path, so history and heading disappear. Static systems routinely commit both at once. They compress a life into a score, and then they photograph the score.

Measuring the point more precisely repairs neither loss. Five decimal places on a single reading still cannot tell the deteriorating person from the recovering one. If someone's name is Finn, a more careful enunciation of "Flynn" does not help. A snapshot may be accurate and still be insufficient.

We prefer to be named than numbered.

Restore time and the shared point becomes one moment on two paths. Administrative categories organize the representation. They do not constitute the person. That is the standing assumption of everything that follows. Life happens in time. Health, money, families, careers — and the institutions that touch them — unfold on clocks that do not necessarily agree. We measure whether people are inside an acceptable region. We rarely ask whether they can stay there. A person can pass the first test and fail the second.

The Demonstration

Four people, one reading.

The room does not open with a framing statement. It opens with these four.

CASE	A SNAPSHOT TOOL REPORTS	WHAT THE HISTORY SHOWS	AN HONEST INSTRUMENT REPORTS
A	Prediabetes	Six years of readings, flat within measurement noise.	Stable. Narrow uncertainty.
B	Prediabetes	Eighteen months, rising steadily through the range.	Deteriorating. Narrow uncertainty.
C	Prediabetes	Three years, falling from a materially higher peak.	Recovering. Narrow uncertainty.
D	Prediabetes	One reading. No prior history of any kind.	Unknown. Classification withheld. Widest uncertainty in the set.

Read the last row again, because it is the expensive one. The honest instrument's best available answer for D is refusal. Not "stable." Not "borderline." Unknown, with the widest uncertainty of the four, and the single observation it did make presented as exactly that: one observation. A snapshot tool prints the same word for all four. An honest instrument names three headings and goes quiet on the fourth. That is the demonstration, and it is the spine of this paper.

The first task is to look beyond appearances without pretending to see what is not there.

Unknown is an acceptable value. A blank with an explanation is better than invented precision. The first duty of a good instrument is not to know everything. It is to know what it knows, and what it does not. There is dignity in that restraint. The four cases are constructed. Nothing about them establishes how often real instruments fail this way, and this paper does not claim to know. What they establish is a standard, and a standard is a thing you can be held to.

Trajectory is not automatically superior to status.

Sometimes the snapshot is exactly what the decision requires. A pilot needs to know altitude now; a college student on a road trip needs to know how much fuel is in the tank now; a clinician needs the blood pressure in front of her now; a retirement committee may need to know whether income replacement is adequate today.

History becomes worth its cost only when the decision depends on what the point cannot tell us. An aircraft descending through 10,000 feet is not the same object as one climbing through 10,000 feet. A car with half a tank 40 miles from home is in a different predicament from the same car 400 miles from home. The same health status means something different when housing, food access, transportation, work conditions, and family obligations are improving than when they are deteriorating.

The same 70-percent income-replacement ratio means something different for a worker accelerating savings at 45 than for one decumulating at 67. Trajectory therefore carries no automatic privilege. It earns its place only when history and heading change what a responsible person would do, and when that additional decision value justifies the added complexity, cost, noise, latency, governance burden, and privacy exposure required to see it.

Two Engineering Rules

What the instrument owes the person.

First. Nothing skips the chain.

observation → state estimate → heading → assessment

A laboratory value is not a state. A state is not a heading. A heading is not a decision. Each arrow is a separate claim, and each claim can fail on its own terms. Nothing should skip that chain and land as a score.

Second. The instrument must know when it should not be trusted.

Every output carries how well-observed it is. We already live this in a car: fuel, speed, miles to empty. A gauge that cannot admit it is broken is worse than no gauge at all.

The non-inference boundary.

On a new highway in Brazil, the satellite navigation put the car in a field. The driver could see the road through the windscreen. The instrument could not see that it could not see.

That is the failure this rule exists for. When the observations cannot support a heading, the instrument withholds the heading and keeps what it did see. It says unknown, not stable. If the evidence does not support a direction, the honest answer is not a more impressive guess. It is: we don't know yet.

Refusal is a limit on the claim, not a limit on care. A clinician may still treat. A plan administrator may still have to decide on time, and usually does. What the box does not get to do is invent the road in order to sound useful.

One distinction matters here and is easy to blur. Not permitted to look is not the same as looked, and saw too little. Uncertainty does not grant a fishing licence.

Don't narrate a road you have not been down.

The Test

What would have to be true.

A standard is worth nothing until someone says how it could fail. This page says how.

If the four-case procedure is ever run as code, everything is written down before the cases are evaluated: what counts as a pass, what counts as too little history, and what must not change the answer. A missing declaration stops the run. A condition changed after seeing the cases voids it.

Then three questions are put to the instrument. Failing any one of them fails the test.

The four again:

A - six years flat

B - eighteen months rising

C - three years falling from a higher peak

D - one reading, no history

HYPOTHESIS	PREDICTED RESULT	WHAT WOULD FALSIFY IT
H1 - Discrimination <i>Can it tell A, B and C apart?</i>	A, B and C receive three distinct classifications matching constructed truth.	Any case misclassified, or any two receiving the same class.
H2 - Uncertainty tracks adequacy <i>Does it admit what it cannot see?</i>	D is refused, and its uncertainty is the widest in the set and does not shrink as the history gets thinner.	D receives a classification, or its uncertainty is not the widest.
H3 - Invariance <i>Does the answer stay put?</i>	Outputs are unchanged when units, date origin, input order or identifiers change.	Any of those transformations changes a classification or a bound beyond tolerance.

A miss on any of the three fails the test. H2 is the one that matters most, because confident output on inadequate data is the failure mode the whole construction exists to expose.

IF IT PASSES, WHAT HAVE WE EARNED THE RIGHT TO SAY?

Very little, and that is the point. A pass would earn one sentence: on these constructed histories, this frozen version of the instrument did what it said it would — it told three apart, refused the fourth, and gave the same answer when we changed the units, the dates and the names.

It would earn nothing about the world. Not a claim about real patients. Not that the confidence figures are calibrated. Not that the instrument predicts anything. Not how often other instruments fail this way. Not that it would survive a working clinic. Not that anybody got better.

The distance between those two paragraphs is not modesty.
It is the paper's own standard, turned on the paper's own claim.

It has not run. Until it does, what you are reading is a standard, not a finding.

Rendering Integrity

A score is not a measurement because it is formatted as one.

A beautiful table can make an inference look measured. An elegant dashboard can make a forecast look observed. A polished score can make a hypothesis look settled. Presentation does not confer knowledge-status.

Every transition from reality to representation must preserve the knowledge-status of what crossed the boundary.

A patient's report should not silently become a diagnosis. A diagnosis should not silently become a claim. A claim should not silently become a health state. An inference should not silently become a fact. A risk score should not silently become a prognosis. A cost category should not silently become low human value. The obligation continues when representation crosses into action. An estimate can inform a rule. A rule can shape a decision. And a decision can alter the conditions under which a person then has to act.

Existing provenance standards - W3C PROV, HL7 FHIR Provenance - record who did what to which record, and when. They do that well, and this argument does not compete with them. They do not carry confidence as a primitive, do not treat disagreement about the same person as a first-class object, and do not record which meaning-operation occurred. The layer required here sits above those standards, not against them.

The flight analog.

A primary flight display does not collapse airspeed, altitude, heading and fuel into a single safety score. It keeps the instruments distinct, and it tells the pilot when a reading is unreliable. Time-to-boundary - how long until the aircraft leaves the acceptable operating region - is a trajectory object, not a status object. A fuel gauge reports what is in the tank now. Miles-to-empty reports how long that lasts.

The first is a status; the second is a trajectory. An instrument that returns a single decimal with no uncertainty flag is asking a person to fly on a collapsed gauge.

Seven Futura Tracks, One Question

Wherever something moves and is only partly seen.

Ask the same question wherever a system is dynamic, path-dependent, constrained, partially observed, and subject to intervention. That describes all seven Futura tracks. The table below invites you to test the analogy in your own track. It is not evidence that the failure is common in any of them.

TRACK	TREATED AS STATIC	WHAT IS ACTUALLY MOVING
AgTech & Sustainability	Yield, input cost, compliance	Soil, water, farm solvency, ecological capacity
Culture, Media & Entertainment	Reach, revenue, hits	Audience direction, platform dependence, franchise coherence
Enterprise	Compliant or not; legacy or modern	Resilience, technical debt, control velocity
Infrastructure & Logistics	Asset up or down; on time or late	Degradation, congestion, disruption-to-recovery path
Innovative Technologies	Accurate or inaccurate model	Calibration, drift, attack surface, confidence collapse
MedTech & Life Sciences	In range or out of range	Physiology heading, trial response, chronic drift
Transport & Aerospace	On course or off course	Decision space expanding or collapsing; fleet and fuel transition

The panel will not tour all seven. It takes one object from MedTech and treats it with the discipline aviation already expects. The table is here so that someone in AgTech, Enterprise or Transport can decide whether the object is theirs and walk over.

Field is what surrounds us.
Trajectory is how we are moving through it.
Navigation is what we do next.

- Where are we?
- Where are we going?
- What is the best route?
- What changed?
- What correction?

The Institutional Version

Health on one clock, wealth on another.

The instrument problem has an institutional twin, and it surfaces at employer renewal as a single question: what is happening to these people, and what may we lawfully do about it?

United States. Health-and-welfare and retirement sit in separate plans, on separate clocks, under separate committees, with separate vendors, funding, testing, filing logic and employee mental models. Welfare runs on the plan year. Retirement runs career-to-n-years. A package can therefore be architecturally two one-axis systems even when it looks comprehensive.

Ireland. A public floor changes the surrounding architecture. The state already occupies part of both axes. Private medical insurance buys access and waiting time, not an accumulating health account. My Future Fund auto-enrolled eligible employees from 1 January 2026, with access at 66. An employer may top up both axes without ever coupling them.

Neither country is simply better. They allocate responsibility differently, and an Irish company entering the United States inherits duties that are less visible under a public floor.

That separateness is real, and it is not an argument. Statutory plan categories, separate committees and separate vendors are facts about how institutions are chartered. They are facts about jurisdiction, not facts about the household living on both clocks at once.

The org chart is not the ontology.

Nor is the mission statement the purpose. What an institution says it exists to do is a claim. What it repeatedly does is evidence of what it is organized to produce. An administrator who says health status and retirement readiness cannot share a plan year is describing the law correctly. That same administrator is not thereby entitled to conclude that the person's predicament is two separate problems.

Representation is not commensuration.

Putting health and financial position in the same representational field does not establish a conversion rate between them. A coordinate system can preserve two distinct dimensions without pretending that one unit of health equals some quantity of money. Aggregation is a further claim again: combining the dimensions into one score would require weights this paper has not established.

Methodology may refuse commensuration; life does not get to refuse the predicament. A household still has to choose among care, cash flow, work, savings, disability and retirement on clocks that interact whether or not the analyst has found a common unit. The institutional question is therefore not whether health and money are the same thing. They are not. It is whether a design can represent their interaction without erasing the distinction.

Two Dials, No Exchange Rate

A thermometer reads 31°C. A hygrometer reads 80%. Both hang on the same wall, and both bear on the same question: can a person work outside this afternoon?

Nothing converts them. There is no rate at which degrees buy humidity. No one has ever paid four points of moisture for a degree.

The heat index does combine them, and it is a real instrument. But it is a further claim: it carries weights somebody had to derive, it holds only across a stated range, and outside that range it is quietly wrong. **Same wall, representation. No exchange rate, commensuration refused. One number, earned: aggregation, and it has a price.**

A coordinate system can preserve two distinct dimensions without pretending that one unit of health equals some quantity of money. Yet, we know that one can convert one hour of labor into a paycheck that can pay health inputs: a solid meal (nutrition) and a bottle of water (hydration). One can sleep an extra hour which gains an essential health input (rest) in exchange for labor (money).

So a Joint Claim has a Bar

Count the dimensions before judging the claim. A proposal that changes health alone is not a health-and-finance design, and naming two variables does not make them joint. A programme is joint only when it can answer four questions in participant documents for the same employee class and plan year:

Which person-level health variable moves? Which person-level financial variable moves? On which shared or explicitly coupled clocks? What is the miss rule, the result that would count as failure of the joint claim?

Enrollment, coverage, sponsor cost, engagement, vendor service levels and a shared open-enrollment week do not satisfy it. If the reader has to infer the miss rule, it is absent.

Whether such programmes are uncommon is an open document question. It is not a premise of this paper.

Governance

Who may steer, and who may not.

An instrument can describe a trajectory. It cannot supply the destination.

Integration must not mean unification of authority. The member is the human referent, the real person the data is about. The longitudinal profile is the informational centre. The group health plan is the financing and fiduciary centre. Those are three different things, and they do not merge merely because one entity pays for all of them.

A physician should not become subordinate to an employer's wish to reduce claims. Plan fiduciaries must be able to see, analyse, administer and pay claims; that duty does not convert the employment function into a clinical office, and it does not let a supervisor inherit the claims file. An analytics platform should not silently become a clinician.

Custody, access authority, use authority,
decision authority and accountability
are five different facts.

Individual trajectory output may flow to a named plan-administration function behind a firewall that can be described. It never flows to the employment function.

Where roles are fused and the firewall cannot be described, it does not flow at all.

And a refusal grounded in authority must be labelled as such, never recoded as insufficient observation.

Responsibility should sit where an actor can actually perceive and influence the relevant part of the path. The org chart records assigned authority. The trajectory records effect. An actor may possess little formal authority and still materially alter the path; another may possess considerable authority and exercise almost none. Governance therefore has to examine not only who is authorized to act, but what the system's actions repeatedly produce.

Several legitimate centres, not one hub. Distributed responsibility is not abandoned responsibility. The institutional counterpart is equally plain: three conditions help people move well inside a system.

ORIENTATION - I can understand the system.

BELONGING - I can contribute without forfeiting my place in it.

CONTINUITY - The system can learn from what we do.

Orientation · Belonging · Continuity
Francesca Pagani, from research by Sofia Menesello.

The Capture Test

Ask it of every entity in the programme, including ours.

Is the model serving better human orientation, or are people being made legible for the model?
Are practitioners serving capability, or enrolling people because a programme exists?
Is administration serving participants, or are participants learning to serve administrative machinery?
Is advice serving the client's intended life, or reframing that life around the institution's preferred services?
Is the venture architecture serving problems encountered, or recruiting problems to populate the architecture?

Purpose is not established by declaration. A system that says it serves people must also be judged by what it repeatedly causes, what it rewards, and the trajectories its operation actually produces. That test applies to us too.

*Before an opportunity enters the architecture, name its status:
Altravia, exploration, candidate venture, or interesting possibility.*

Cost is a constraint. It is not the purpose. A system that reduces cost by reducing beneficial care, shifting burden, or improving risk selection without improving function has not stewarded the asset it named.

The Person, Not The Mark.

An duine, ní an marc.

A score can be useful.
A category can be useful.
A model can be useful.
But none of them gets the last word.
The representation should conform to the person.
The person shouldn't conform to the representation.

People are named, not numbered.

Status

What is settled, and what is not.

CLAIM	STATUS	WHAT WOULD SETTLE IT
Trajectory is the right unit of analysis	Proposed	No measurement settles a choice of unit — only a more useful account of the same evidence.
The instrument distinguishes trajectory from a point reading	Untested	A predeclared conformance test on constructed cases meeting the stated discrimination criterion.
Reported uncertainty tracks observation adequacy	Untested	The same predeclared conformance test showing that inadequate cases are refused and carry the greatest reported uncertainty.
The uncertainty output is statistically calibrated	Untested	A prospective cycle testing interval and event-probability calibration on held-out observations.
A heading predicts rather than describes	Untested	A predeclared prospective study against last-point and simple-trend comparators, after the necessary partner and ethics conditions are met.
Intervention changes a trajectory	No evidence	A controlled, adequately powered study that may not begin until the prediction study exists.
Joint health-and-finance employer programmes are uncommon	Untested	A predeclared document sample against the four-part joint-claim test, with sample counts and falsification criteria fixed before coding begins.
Health and finance form a navigable plane	Untested	A fixed-horizon check showing joint out-of-sample value beyond separate margins.

Three states only.

Proposed is articulated in prose in this paper. People encounter changing conditions and make consequential decisions whether or not a validated methodology yet exists for representing those conditions. Methodological uncertainty limits what an instrument may claim; it does not suspend the reality a person must navigate.

Untested means a named test could settle it and that test has not run.

No evidence means no study exists, and none may begin until an earlier one does.

The one that moves next.

Six empirical programmes could follow from this paper.
Most require additional clinical, ethical, or longitudinal-data conditions before they can begin.

One does not. The employer-programme document test can be run against ordinary plan materials - summary plan descriptions, wrap documents, open-enrollment guides, Form 5500 filings - without waiting for a clinical study. Its sample and falsification criteria must be fixed before coding begins.

That is the next empirical test we intend to run.

The Ask

What the forty-five minutes are for.

The panel is not a thesis presentation. The live question is narrower and more useful than the paper: how much uncertainty can a commercial instrument expose before it becomes unusable?

The rival is not abstract. It is the objects a renewal meeting already has on the table - enrollment, unit cost, trend, risk tier, stop-loss - measured against whatever a trajectory estimate would actually add. Several answers are defensible. This paper can name the disagreement. Only the room can have it.

Four people. One reading. One refused classification. And one question: what would a defensible instrument owe the person it measures, and what does it cost to build one that will go quiet?

What We Are Looking For

Partners willing to take one specific population, one specific set of documents, and one specific decision seriously enough to stay in the room until something testable comes out of it: governments, institutions, employers, insurers, clinicians, civic leaders.

Our instrument for that is a design charrette: one structured day, the right people, mapping what is real and producing a plan coherent enough to act on. Not a report. A channel.

One condition, and it is ours rather than yours: the charrette names its own miss rule before it begins. We write down, in advance, what result would count as the day having failed. A process that cannot fail has not been designed. It has been staged.

Find us in Limerick. Or write first, and we will arrive with your documents already read.

Ní staid atá le riaradh é bláthú an duine. Is conair é atá le stiúradh i gcaitheamh an ama. Féadfaidh beirt a bheith ar an bpointe céanna gan a bheith ar an turas céanna. Is fearr linn ainm a bheith orainn ná uimhir a bheith orainn. Nuair nach leor an fhianaise chun treo a dhearbhu, ba cheart don ghléas cloí leis an méid atá feicthe aige agus a admháil nach eol dó cén treo atá i gceist.

Human flourishing is not a state to be administered. It is a path to be steered through time. Two people may stand at the same point and not be on the same journey. We would rather bear a name than a number. When the evidence is not enough to confirm a heading, the instrument should keep to what it has seen and admit that it does not know the heading.

Da dove viene ALTRAVIA?

Altravia è figlia dell'amore per ciò che esiste già, ma che non sempre siamo capaci di vedere.

Nasce dal desiderio di guardare con occhi diversi: fermarsi, osservare, riconoscere valore in ciò che il tempo, l'abitudine o la distanza hanno reso invisibile. Perché a volte ciò che manca non è qualcosa di nuovo da creare, ma qualcuno disposto, finalmente, a vedere ciò che è già lì e a prendersene cura.

Il lato italiano di Altravia porta con sé questo sguardo sul tempo. Siamo cresciuti in luoghi in cui epoche diverse convivono, dove *ciò che viene prima non scompare, ma continua a essere abitato, trasformato e reinterpretato*. Ogni cambiamento nasce dentro una storia già iniziata, fatta di conoscenze, relazioni, tentativi, errori e intuizioni che continuano a vivere nel presente.

Per questo, innovare non significa necessariamente sradicare ciò che esiste per sostituirlo con qualcosa di nuovo. Può significare riconoscerne le radici, dare voce e valore a ciò che è stato trascurato, *riparare le connessioni* che nel tempo si sono interrotte e nutrirlle di nuove idee.

Ma la cura non ha come unico scopo la conservazione. Una radice curata vuole anche generare. Altravia crede nella fecondità senza timidezza: nell'abbondanza di idee, relazioni, imprese, opere e possibilità che può nascere quando qualcosa di vivo trova finalmente le condizioni per esprimere tutto ciò che porta dentro di sé.



STRADE ANTICHE DESTINAZIONI NUOVE

Altravia nasce forse proprio qui: nell'incontro tra uno sguardo capace di vedere ciò che già esiste, la cura necessaria per riparare ciò che si è interrotto e il desiderio di aiutare ciò che è vivo a diventare generativo. Perché non custodiamo un frutteto soltanto perché sopravviva.

Lo custodiamo perché fiorisca, dia ombra, nutra, produca frutti e semi e perché da quei semi possano nascere frutteti che oggi non siamo ancora capaci di immaginare.

Sofia Menesello · Francesca Pagani



1218 Webster Street · Houston
Texas 77002 · United States

altravia.studio