

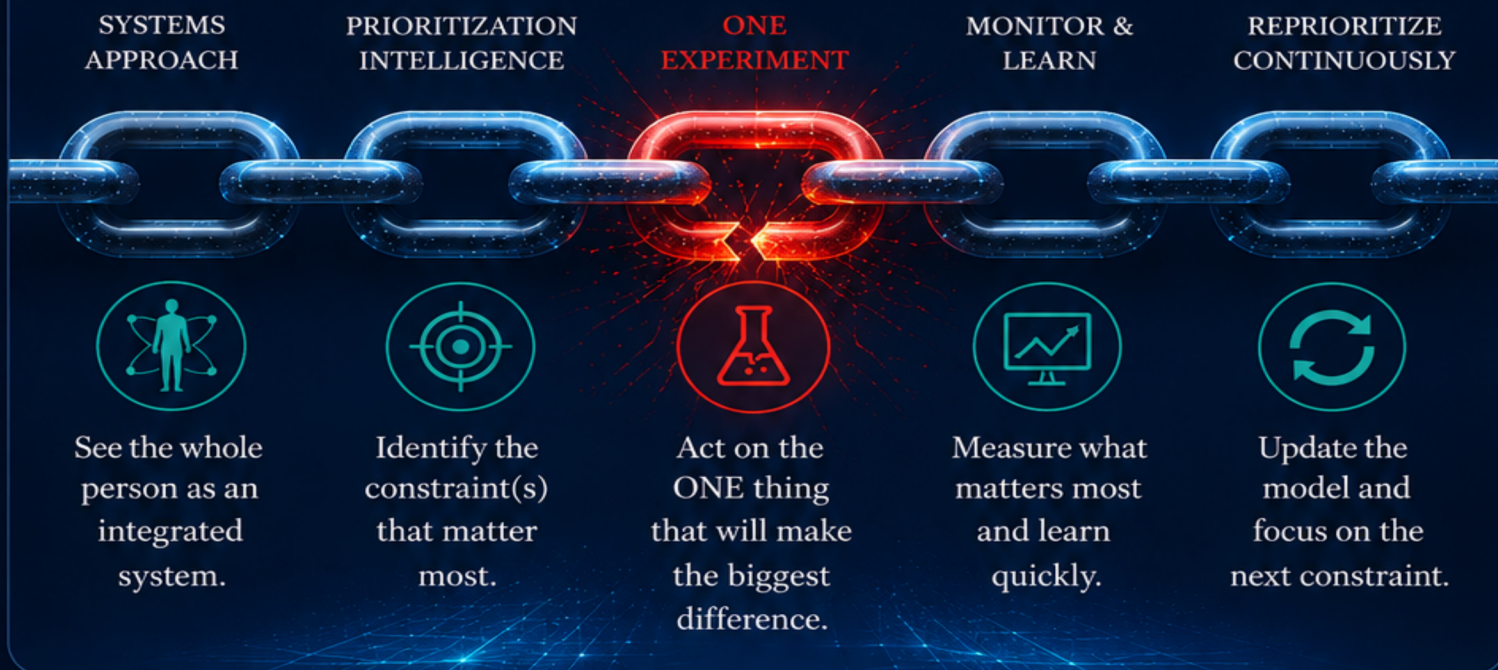
CONSTRAINT-FOCUSED P5 MEDICINE

Applying Theory of Constraints,
AI and Digital Twins to Prioritize
Healthspan Interventions

Dr. Alan Barnard | Dr. Graham Simpson



ONE WEAKEST LINK LIMITS YOUR HEALTHSPAN THE MOST



Eternity
HEALTH AI

DECISION SCIENCE FOR 100+ YEARS OF HEALTHSPAN

Abstract

Background. A century of progress in vaccination, antimicrobial therapy, surgery, imaging, intensive care and molecular biology has transformed acute disease and extended human lifespan. The central task of medicine has shifted accordingly — from helping people survive disease to helping them stay healthy across increasingly long lives.

The problem. Systems biology, molecular diagnostics, wearable sensors, artificial intelligence (AI) and Digital Twins now let clinicians observe human physiology in unprecedented breadth. This abundance has created a new and largely unaddressed difficulty: for any individual, hundreds of measurements may each be relevant, interacting and abnormal, yet clinical time, patient capacity and resources are finite. Our ability to *identify* what is abnormal has advanced far faster than our ability to *prioritize* what to address first.

Proposal. We propose **Constraint-Focused P5 Medicine**, a conceptual framework that extends Leroy Hood's P4 Medicine (Predictive, Preventive, Personalized, Participatory) with a fifth principle — **Prioritized** — and integrates systems biology with the Theory of Constraints (TOC), AI, Digital Twins and Decision Intelligence.

Its central claim is that, at any moment, an individual's healthspan is disproportionately limited by a small number of **physiological constraints**, and that identifying and relieving the dominant constraint first may produce broader system-wide benefit than treating many abnormalities in parallel.

We introduce four constructs — **Constraint Physiology** (the study of limiting physiological constraints), **Constraint Intelligence** (the disciplined identification of the dominant constraint), the **Energy Constraint Hypothesis** (a leading, falsifiable candidate constraint), and **Decision Intelligence** (resolving the *behavioral* constraint that limits implementation).

Key testable hypothesis. In a defined population, prioritizing interventions on an individually identified dominant physiological constraint will produce greater improvement in validated healthspan outcomes, at equal or lower burden, than comprehensive non-prioritized care.

Implications. The framework reframes the clinician's highest-value contribution from measuring more to determining what matters most, and it makes prioritization an explicit, testable object of research. It is offered as an evolving, falsifiable framework — an invitation to investigation, not a claim of established doctrine.

Keywords: healthspan; systems medicine; Theory of Constraints; prioritization; P4/P5 Medicine; mitochondrial energy metabolism; Digital Twins; decision science; personalized medicine.

Table of Contents

ABSTRACT.....	1
1. INTRODUCTION.....	3
2. THE EVOLUTION OF MEDICINE	4
3. THE CONVERGENCE	5
4. THE PRIORITIZATION PROBLEM.....	6
5. CONSTRAINT PHYSIOLOGY	7
6. CONSTRAINT INTELLIGENCE.....	12
7. FROM THE TRIVIAL MANY TO THE VITAL FEW AND CRITICAL ONE	16
8. THE ENERGY CONSTRAINT HYPOTHESIS	17
9. DECISION INTELLIGENCE: FROM KNOWING TO DOING	22
10. HOW CONSTRAINT-FOCUSED P5 MEDICINE DIFFERS.....	25
11. TESTABLE PREDICTIONS	26
12. ANTICIPATED OBJECTIONS AND RESPONSES	27
13. LIMITATIONS.....	28
14. IMPLICATIONS FOR CLINICAL PRACTICE	28
15. MONITORING THE ONE EXPERIMENT	30
16. A RESEARCH AGENDA.....	31
16. CONCLUSION.....	32
EPILOGUE: STANDING ON THE SHOULDERS OF GIANTS	33
APPENDIX A: WORKING DEFINITIONS	34
APPENDIX B: A CLINICAL WORKFLOW	35
DISCLOSURE	36
REFERENCES	36

1. Introduction

Medicine advances through a recurring pattern: a new way of observing the body enables a better question, and the better question reorganizes practice. New knowledge rarely replaces what came before; it extends it, revealing relationships earlier generations could not see. This paper is written in that tradition. It does not seek to displace existing medicine but to add one capability the current moment increasingly demands — a disciplined way to decide *where to begin*.

We believe the best innovations begin with the right questions, and the Theory of Constraints is, above all, a discipline for asking them. Applied to a human life, it asks three in sequence.

First, what is the ultimate goal of the system? Here it is a healthspan of a hundred years or more, lived with function, independence and meaning. That goal is not merely survival and growth; at every moment it also requires the best possible allocation of the body's scarce daily resources.

Second, what are the few performance dimensions that must be measured to judge progress toward that goal, to reveal the weak links and the single weakest link, and to surface the risk factors that could derail it?

Third, what is currently limiting the system: its capacity, capability and availability to meet the demands of daily life, and the durability of each subsystem to keep meeting them across a hundred years and beyond? The remainder of this paper is, in effect, a disciplined attempt to answer these three questions.

These questions must be asked of every major system, and the first is really several.

Capacity asks whether a system can meet the demand placed on it today; capability asks whether it can do so in the right way, not merely in bulk; and availability asks whether it can do so when and where it is needed.

Durability then asks something different: on its current trajectory, is the system likely to remain sufficient to support a healthy life beyond a hundred years?

A system can satisfy the first three today and still fail the fourth, much as a machine that runs well now may wear out before the end of a century-long production run. Two decision thresholds give these judgments clinical meaning. Crossing a healthspan threshold signals a likely loss of function, resilience or independence; crossing a lifespan threshold signals a material rise in the risk of premature mortality. Some systems influence healthspan without carrying a direct lifespan threshold, and are managed accordingly.

The success of twentieth-century medicine created a twenty-first-century problem. As people survive once-fatal illnesses and live for decades with several interacting chronic conditions, the defining question is no longer only how to extend life but how to preserve health, function and independence across that longer life.

At the same time, the quantity of biological information available for any one person has grown by orders of magnitude.

A single assessment may combine genomic, proteomic and metabolomic data, hundreds of laboratory analytes, imaging, body composition, and continuous wearable streams describing sleep,

glucose, heart-rate variability and activity. AI surfaces still more associations; Digital Twins promise individualized simulation.

Paradoxically, more information does not simplify the clinical decision — it complicates it. As measurable variables multiply, so do candidate explanations, competing interventions and plausible priorities.

In healthspan medicine this is the rule rather than the exception: a typical middle-aged patient may simultaneously show early insulin resistance, elevated inflammatory markers, suboptimal micronutrient status, declining muscle mass, disrupted sleep, reduced heart-rate variability and chronic stress.

Each finding may deserve attention. Not all can be pursued at once, and treating each in isolation offers no guarantee of the greatest overall benefit.

Contemporary preventive medicine has become extraordinarily good at *identification* — of risk factors, biomarkers and interventions — and comparatively underdeveloped at *prioritization*: the systematic determination of the order in which interventions should be undertaken for a given individual. This paper asks whether medicine would now benefit from an additional systems perspective built specifically for clinical prioritization, and proposes one.

2. The Evolution of Medicine

Each era of medicine has been defined not only by new treatments but by a new way of understanding health. In the nineteenth century, Claude Bernard's *milieu intérieur* shifted attention from individual organs to the regulation of the body's internal environment; Walter Cannon later formalized this as homeostasis.

Both established that health is not a static state but continuous adaptation. In the twentieth century, advances in pathology, microbiology, pharmacology, surgery and imaging transformed clinical care, while Sir William Osler reminded physicians to treat the patient rather than only the disease — an enduring counterweight to technological sophistication.

Over the past three decades, systems biology reframed physiology again, seeking to explain how interactions among genes, proteins, cells, organs and environment collectively determine behavior (Kitano, 2002).

Dr. Leroy Hood then reframed that ambition into a program for medicine itself. With Weston he argued for care that would become predictive, preventative and personalized (Weston and Hood, 2004), and in the same year, with Heath, Phelps and Lin, showed how systems biology and new technologies could make prediction and prevention practical (Hood, Heath, Phelps and Lin, 2004). Adding participation as a fourth principle completed the formulation now known as P4 Medicine. As Hood and Friend put it, "*Medicine will move from a reactive to a proactive discipline over the next decade—a discipline that is predictive, personalized, preventive and participatory (P4).*" (Hood and Friend, 2011). Dr. Nathan Price, working with Hood, helped turn that vision into evidence and practice: their study of 108 individuals demonstrated that dense, longitudinal "data clouds" reveal health trajectories rather than isolated snapshots (Price, Magis, Earls et al., 2017), and together they set out the case for scientific wellness in *The Age of Scientific Wellness* (Hood and Price, 2023). Our proposal begins where theirs arrives.

The fourth P has an origin worth recording. By 2002, Dr. Hood was describing P3 medicine: predictive, preventive and personalized. The fourth principle arrived from outside medicine, in conversation with Google co-founder Larry Page. In Dr. Hood's own recollection: "*Larry said people participating in their own care is just as important as the other three Ps, and I immediately realized that's true.*" (Hood, quoted in PMC2957728). That observation added the principle that makes the other three usable, because a plan the patient does not own is a plan that does not happen. We record it for two reasons. It credits an insight that changed the shape of the field, and it is itself an instance of the pattern this paper is concerned with: a limiting factor named by someone viewing the system from a different discipline. Our own attention to the behavioral constraint on implementation (Section 8) is a direct descendant of that fourth P.

Dr. David Furman and others deepened the understanding of chronic inflammation and immune aging using AI and Digital Twins as central drivers of analyzing and identifying ways to increase healthspan.

A parallel current runs through Dr. Graham Simpson (MD)'s five decades of clinical practice. In *Era III Medicine* (Simpson, 2025), he argues that medicine is entering a phase defined by the integration of systems biology, personalized diagnostics and continuous optimization, and that chronic disease is better understood as the behavior of an integrated metabolic system than as a collection of isolated disorders — with cellular energy metabolism increasingly at the center of aging and chronic disease.

While medicine was integrating, a separate discipline was learning how complex systems improve. Through the **Theory of Constraints**, Dr. Eliyahu Goldratt showed that the performance of any complex system is governed by a small number of limiting constraints, and that improving the dominant constraint typically yields disproportionately large gains — whereas optimizing everything at once does not.

Dr. Alan Barnard (PhD) has spent three decades extending TOC into decision science, leadership and human performance. In *From Many to Few to One* (Barnard, 2026) he sets out the **ONE Thing Focusing Cycle** — a repeating discipline that moves from the *one goal* that matters most, to the *one constraint* that most limits it, to the *one problem* at its root, to the *one conflict* that keeps that problem unresolved, to the *one innovation* that breaks the conflict, to the *one experiment* that tests it — then repeats on the next constraint. It is a practical method for sustaining focus by solving the most limiting constraint, then the next, one at a time.

Taken separately, each contribution advanced one facet of health. Taken together they suggest something larger. Systems biology explained biological complexity; P4 Medicine redefined the goal; diagnostics, wearables and AI expanded observation; the Theory of Constraints offers a disciplined method for deciding *where to act*; and decision science offers methods for turning insight into sustained action. It is the **convergence** of these streams that motivates the framework proposed here.

3. The Convergence

Scientific change rarely follows a single discovery. More often, independent advances mature at the same time and create possibilities none could produce alone. Six such capabilities are now converging in medicine: systems biology; high-resolution diagnostics and multi-omics; continuous

physiological monitoring; artificial intelligence; Digital Twins; and decision science. Individually each is transformative. Together they make it feasible, for the first time, to observe an individual's physiology continuously, model it computationally, and reason about which part of the system most limits the whole.

Crucially, convergence does not by itself solve the clinician's problem — it intensifies it. Each new capability adds measurements, associations and options. The missing layer is not another sensor or another model but a disciplined method for turning abundant observation into a small number of high-value actions. That is the gap Constraint-Focused P5 Medicine is designed to fill.

4. The Prioritization Problem

For most of medical history, physicians were limited by *scarcity* of information; diagnosis depended on recognizing disease once it had become clinically apparent. Today's clinician faces the opposite condition — *abundance*. A comprehensive assessment may span genomics, advanced laboratory panels, inflammatory and metabolic markers, endocrine profiles, microbiome analysis, body composition, cardiovascular measures and continuous wearable data. This yields an unprecedented picture of the individual — and a decision problem that grows with every added variable.

Three features make the problem hard. First, patients rarely present with one dominant abnormality; they present with many, interacting in nonlinear and often poorly understood ways. Second, resources are finite — clinical time, and equally the patient's attention, finances and emotional bandwidth. Third, and least appreciated, treating every abnormality in parallel is not obviously optimal even when it is possible: in a system of interacting parts, diffuse effort can dissipate across many small improvements while the factor most limiting the whole remains unaddressed.




Figure 1. From P4 to P5. The evolution from Hood's four principles (Predictive, Preventive, Personalized, Participatory) to five, adding Prioritized via the Theory of Constraints.

A familiar principle from systems science applies.

A chain is only as strong as its weakest link.


Complex adaptive systems seldom improve by optimizing every component simultaneously; they improve fastest when the current limiting factor is identified and relieved, which often changes the behavior of the entire system and reveals the next opportunity. Whether — and how far — this principle transfers from engineered systems to human physiology is an *empirical* question. We take it seriously precisely because it is unproven, and because, if even partly true, it would reorganize how healthspan medicine is practiced.

Framed this way, the chain does something most methods do not: it gives three answers, not one — and the third is the one almost everyone misses.

- **Many** of a person's systems are already strong enough; effort spent there is pure waste of time, money and the patient's finite willingness to keep going.
- A **Few** are genuinely weak; strengthening them creates potential but no result, because the chain still fails at its weakest point.
- And there is always **One** weakest link; only here does effort convert into system-wide improvement — and only until it is no longer the weakest.

From many, to few, to one — and then again. The move from *few* to *one* is where reductionism does its quiet damage: correcting a real but non-limiting abnormality makes a number improve while the patient does not, and both clinician and patient acquire direct personal evidence that “this does not work.”

A Chain Is Only as Strong as Its Weakest Link



Strengthening a strong link adds nothing. Only the weakest link limits the whole.

Figure 2. A chain is only as strong as its weakest link. Many strong systems, a few weak, and one weakest link that limits the whole — the constraint identified first.

That question — which of the few weak links is *the* weakest — leads directly to the framework's foundational construct: **Constraint Physiology**.

5. Constraint Physiology

The Theory of Constraints holds that a complex system's performance is limited by a few constraints that set its ceiling under current conditions; relieve the dominant one and another becomes limiting, so improvement is a continuous cycle of identifying and relieving successive constraints rather than optimizing all components at once.

Human biology shares many properties of complex adaptive systems: continuous feedback, nonlinear interactions, and outcomes of the whole that cannot be predicted from the parts in

isolation. This resemblance does not license importing manufacturing logic wholesale — biological systems are far more complex, redundant and adaptive than factories, and the analogy must be held with humility. But the resemblance is strong enough to justify serious exploration.

The central proposition of Constraint Physiology is deliberately modest:

At any given time, certain physiological processes exert disproportionately greater influence on an individual's overall health than others. Identifying and relieving the dominant constraint may therefore produce broader, system-wide improvement than treating each abnormality independently.

Constraint Physiology does not replace physiology, pathology or diagnosis. It adds one question to them. Traditional physiology asks *how do biological systems function?*; pathology asks *how does disease disrupt them?*; Constraint Physiology asks: **which physiological process is currently limiting the overall performance of this person's biology?**

The distinction is subtle but consequential — two patients with similar laboratory findings may have different limiting constraints, and two patients with different diagnoses may share one. Attention shifts from isolated biomarkers to the *hierarchy* of factors that determine overall health.

Three propositions follow, offered as hypotheses rather than established fact:

1. **Non-equivalence.** Not all physiological abnormalities contribute equally to overall health.
2. **Dominance.** At a given time, a small number of constraints (often one) disproportionately limit healthspan.
3. **Dynamism.** Constraints are not static; relieving one allows another to emerge, so prioritization is continuous.

A patient limited chiefly by severe sleep disruption may find that restoring sleep improves insulin sensitivity, inflammation, hormonal regulation, exercise capacity and cognition together — after which a different factor, such as inadequate muscle mass or a nutritional deficit, becomes limiting. Health optimization thus becomes a dynamic process of successive improvement rather than a fixed treatment plan.

The Dynamic Constraint Hierarchy

Relieving the dominant constraint changes the system and lets the next one emerge.



Prioritization is continuous, not one-time.

Figure 3. The dynamic constraint hierarchy. Relieving the dominant constraint changes the system and allows the next constraint to emerge; prioritization is continuous, not one-time.

The value of these propositions is entirely contingent on whether acting on them improves outcomes. That question is open — which is exactly why it merits systematic study rather than assumption.

A necessary refinement: constraint topology. The most common objection to a weakest-link framing is that a patient rarely has one weakest link. The objection is correct, and it is answered not by abandoning the chain but by classifying it. Any complex adaptive system can be characterized by the number of constraints it carries and the interdependence between them, and three topologies prove clinically useful.

Three Types of System, Three Different Strategies

Systems can be classified by how many constraints they have and how those constraints interact.




Figure 4. Three constraint topologies. The number of constraints and the degree of interdependence between them determine which strategy applies.

Type 1 – Simple. Several links may be weak, but one is clearly the weakest, separated from the next by more than measurement error. The strategy is to identify the weakest link and strengthen it.

Type 2 – Complex. Multiple links are weak, and a small number of them – not many – lie close enough together in predicted performance that all qualify as weakest. They influence one another but do not interact directly; primary causes remain local, and the constraint set is relatively stable over time. The strategy is to identify the few weakest links and strengthen them in priority order.

Type 3 – Chaotic. Multiple links are weak, they lie very close together, and the identity of the weakest changes continuously – within days or hours – according to the demand placed on energy and the supply available at the time. Treating individual links is largely futile because the target moves faster than the intervention. The strategy is to move one level up and identify the scarce resource that is under-supplying every link at once.

The business analogy is instructive. In a cash-constrained enterprise, marketing, research, production and procurement are all chronically underfunded; shortcuts are taken and the cheapest option is bought rather than the right one. The enterprise appears to lack capacity, capability, availability and supply. It lacks none of these. It lacks cash, and every apparent local weakness is downstream of that single higher-order constraint. In a patient, the equivalent scarce resource is most often energy (§7).

A worked example: ten systems, one hundred years. Suppose healthspan depends on ten major physiological systems and the patient’s goal is 100+ years of healthy, independent life. For each

system the Digital Twin estimates predicted healthspan – the age at which that system ceases to support full function – and predicted lifespan, the age at which its failure would end life. The difference between them is the expected duration of disease.




Figure 5. The same three topologies expressed as predicted healthspan per system against a 100-year goal. Solid bars show healthspan; hatched extensions show the years of disease that follow.

Simple. Six systems are projected to support 100+ years. Three are weak, with healthspans of approximately 90, 85 and 78 years against lifespans of 92, 87 and 87. One – here the cardiovascular system – sits at 68 years. It is separated from the next weakest by roughly a decade, so the constraint is unambiguous. Note also that this constraint shows almost no gap between healthspan and lifespan: a cardiovascular constraint may take the whole system with very little warning.

Complex. Three systems cluster at approximately 72, 70 and 68 years. Because Digital Twin predictions carry an uncertainty of roughly two to three years, systems separated by four years or less cannot be ranked with confidence. They are therefore treated as three weakest links, all requiring high priority, rather than falsely ordered. These constraints are also relatively stable and will remain limiting for extended periods.

Chaotic. The weak systems are not merely close in predicted performance – the ordering itself changes. Which system is limiting depends on how the body has allocated energy and other scarce resources that week, and on the stressor load those systems have absorbed. A ranking generated on one day may not hold a week later. This instability, rather than the number of weak links, is the defining feature of the chaotic topology and the reason link-level treatment fails.

What prioritization produces. Prioritized does not mean listing everything that could be improved. It means descending from the many possible treatments and dosages, to the few that must change for the system to improve at all, and then to the one to act on next. All three topologies produce an output of the same shape: a prioritized list with the ONE at the top and the remaining high-priority candidates beneath it in likely order. What differs is the source of the ONE – the single weakest link in a simple system, the few weakest in a complex one, the higher-order scarce resource in a chaotic one.

What “Prioritized” Means

From the many, to the few that must change, to the ONE to do next.




Figure 6. Prioritization descends from many, to the few that must change, to the ONE to act on now, communicated in a form the patient can act on and evaluated as a single experiment.

Two properties of that list govern how it is used. First, the list is not executed sequentially: the ONE is implemented, its impact measured, the model updated and the list reprioritized before the next action is selected. Because relieving one constraint changes the system, every item below the ONE is provisional by design. Second, “ONE” denotes one experiment rather than one treatment. It may be a coordinated bundle – the highest-priority change in nutrition, in exercise, in sleep and in stress management, together with a supplement or medication where warranted – initiated together and evaluated as a single unit. We refer to this as the ONE Experiment.

Identifying the constraint and naming the treatment are necessary but not sufficient. Implementation requires that the plan be stated in a form the patient can hold: what to stop doing, what to start doing, at what dosage, for what duration, and what result to expect. Stating the expected result is what converts a prescription into an experiment – it gives physician and patient a shared criterion against which to judge response, makes the feedback loop real rather than impressionistic, and allows dosage or duration to be adjusted on evidence rather than impression. It also reduces the overwhelm that a long undifferentiated problem list reliably produces.

A natural-first bias. Constraint-Focused P5 Medicine carries an explicit bias toward the least invasive effective intervention: what to stop and start in nutrition, exercise, sleep and stress management first; then supplements; then medications; then medical procedures. The rationale is not ideological. The first tier is lower in cost and risk, builds physiological capacity rather than substituting for it, and creates fewer dependencies to unwind later. It is a bias rather than a rule, and it is overridden where a patient is in crisis or at imminent risk, or where delay would itself increase risk.

A necessary refinement: two kinds of constraint. A physiological constraint can limit health in two distinct senses.

One limits the system’s current *capacity* to meet the demands of daily life – a **healthspan constraint**. Another may not limit current function at all, yet still cap **life expectancy**, because a single organ system is quietly declining toward failure – a **lifespan constraint**.

The two need not be the same organ or process. Measuring the functional reserve and *rate of decline* of each major organ system therefore matters even when that system is not currently limiting. Digital Twins are well suited to estimating both: the constraint that most limits current capacity, and the organ system that most limits remaining lifespan.

A candidate taxonomy of physiological constraints.

Constraint Physiology requires a shared vocabulary of *what* can be limiting. We propose the following non-exhaustive, working taxonomy of candidate physiological constraint domains, offered for refinement and empirical testing. Any of these may be dominant in a given individual at a given time:

- (1) **Energy** — mitochondrial capacity; the balance of energy supply and demand.
- (2) **Metabolic regulation** — insulin sensitivity, glycemic control, metabolic flexibility.
- (3) **Inflammation & immune regulation** — chronic low-grade inflammation, immune aging.
- (4) **Structural/musculoskeletal capacity** — muscle mass, strength, physical function.
- (5) **Sleep & circadian regulation** — sleep architecture, recovery, circadian alignment.
- (6) **Endocrine/hormonal regulation** — thyroid, adrenal, sex-hormone and related axes.
- (7) **Nutritional sufficiency** — macro- and micronutrient status, protein adequacy.
- (8) **Cardiovascular-autonomic function** — vascular health, autonomic balance (e.g., HRV).
- (9) **Neurocognitive & psychological resilience** — chronic stress load, mood, cognitive reserve.
- (10) **Gastrointestinal & microbiome function** — barrier integrity, absorption, microbial ecology.

A central research task (see §12) is to determine whether such domains can be measured, ranked and validated as *constraints* — factors whose relief produces disproportionate, system-wide benefit — rather than merely as risk factors.

6. Constraint Intelligence

If Constraint Physiology asks *which* process is limiting health, the immediate follow-on is operational: **how is the dominant constraint identified reliably in an individual?**

Historically this was the province of clinical judgment developed over years of practice. That judgment remains central and is unlikely to be replaced. But the volume and complexity of modern data now exceed what any clinician can integrate consistently without computational support. The challenge is not a deficit of expertise; it is the scale of available knowledge.


We describe the integrated capability needed to meet it as **Constraint Intelligence**: the disciplined process of identifying the physiological constraint currently exerting the greatest influence on an individual's health, by integrating clinical expertise, advanced diagnostics, continuous monitoring, AI and computational modeling. Constraint Intelligence is not a new diagnostic device; it is a **decision-support framework**. Its purpose is not to enumerate abnormalities but to rank them — to determine which one most deserves attention now.

It draws on five complementary inputs:

- (1) clinical medicine — history, examination and physician judgment, which supply context no algorithm can replace;
- (2) advanced diagnostics — objective measures of metabolism, inflammation, endocrine and immune function, nutrition, cardiovascular health and biological aging;
- (3) continuous monitoring — wearables that extend observation into real life across weeks and months;
- (4) artificial intelligence — pattern recognition across many variables and longitudinal data, used to augment rather than replace clinical reasoning; and
- (5) Digital Twins — still emerging, but eventually able to simulate candidate interventions computationally before they are tried clinically.

Constraint Intelligence

Five complementary inputs integrated into one prioritized output.



A decision-support framework under physician oversight, not a diagnostic device.

Figure 7. Constraint Intelligence. Integration of clinical expertise, diagnostics, monitoring, AI and Digital Twins into a single prioritized output — the dominant constraint.

Represented this way, the body is a chain of interacting systems whose overall throughput, healthspan, is limited by its weakest link and reaches its end when the first system crosses its failure threshold. A digital twin of the individual can be used in two complementary ways: analyzing each system considered in isolation, and analyzing all systems considered together, coupled as they actually are, supplying and burdening one another.

The difference between the two reads is itself informative, since it estimates how much is missed when the body is treated as a set of separable parts. Because every such estimate carries uncertainty, the twin is most honest when it expresses the weakest link not as a single verdict but as a probability distributed across candidate constraints, informed by inexpensive functional proxies such as cardiorespiratory fitness and bioelectrical phase angle. We describe the twin here only at the level of principle; its specific implementation is developed and refined separately, and is intended to keep improving with use.

Rather than reduce this to a formula, we are explicit about the problem it solves and the criteria a solution must meet. Every clinician, patient and administrator faces a version of the same challenge,

made sharper each year by the near-exponential growth in available tests, treatments and competing goals: how to separate the trivial many findings that are real but not limiting, from the vital few that genuinely interact to hold health back, and from the critical one — the weakest link whose relief will do the most for the whole.

What matters is that this separation serve the global optimum, the health of the whole system, and not a local one. It is easy to improve a single measure while harming the system: driving one number to its ideal at another's expense, or relieving a constraint that was not in fact limiting.

Prioritization in the fifth P is therefore defined by that objective, and by the criteria a sound choice must satisfy, not by a fixed score.

The criteria are the questions a good decision answers: whether relieving a constraint yields disproportionate, system-wide benefit rather than a local gain; how confident we are that it is genuinely limiting rather than merely abnormal; how feasible its relief is given the patient's biology and circumstances; how urgent and how dangerous inaction is; and how much acting on it would teach us. These are considerations to be weighed together, not terms to be multiplied.

We make a deliberate claim about how that weighing can be done. No fixed heuristic or algorithm can perform it reliably, because a human being is a complex adaptive system that is volatile, uncertain, complex, constrained and ambiguous (VUCCA), in which the effect of relieving any one constraint depends on the state of all the others and shifts as the system adapts. A rule simple enough to write down cannot hold that; a rule complex enough to try would be brittle and impossible to audit.

The only reliable route to a genuinely system-level, global optimization is a personalized digital twin that can represent the interactions, rehearse the alternatives and report its answer with honest uncertainty, used under expert human oversight and improved continuously as data accumulate.

Note on terminology. VUCA, introduced by Bennis and Nanus (Leaders: The Strategies for Taking Charge, 1985), describes four conditions of a fast-changing world: volatility, uncertainty, complexity and ambiguity of competing objectives.

VUCCA is Dr. Alan Barnard's extension, which adds a fifth condition, constraints, on the grounds that the analysis, improvement and management of any complex system depends on what limits it. This fifth element of VUCCA, constraints, is deliberately distinct from the fifth P of Medicine, Prioritized.

Constraints are a condition to be reckoned with; Prioritization is the discipline of choosing which constraint to relieve first. The fifth P we propose is Prioritized, not Constraints.

Here a second principle earns its place, the same one that governs our measurement. As Eli Goldratt put it, *"it is better to be approximately right than precisely wrong."* A digital twin need not model a person perfectly, which is impossible, only be approximately right about the interactions that matter and honest about what it does not know.

We have reached the point where personalized twins are approximately right, and when they are asked the right, system-level questions that the Theory of Constraints supplies, they become accurate and personalized enough to serve as the prioritization mechanism that predictive, preventive and participatory care has until now lacked. The twin need only relieve each constraint to the level of good enough, the point at which it no longer limits the system, before attention moves to the next; perfection in any one part is neither necessary nor safe.

This mirrors how the first four P's became practical at all. Prediction, prevention, personalization and participation were not delivered by clinicians unaided; they became possible through the convergence of abundant data, AI, digital twins and human expert oversight. The fifth P is enabled in exactly the same way. Its single addition is a focusing lens, the Theory of Constraints, which directs that combined capability at the one question that most improves the whole: what is limiting the system now?

A single lead measure for the whole system. Identifying a constraint yields a second, often-overlooked benefit: it gives clinicians one system-level *lead* measure of whether an intervention is actually working. In operations, once the bottleneck is known one does not track every machine — one asks whether the bottleneck is now producing more, being starved less, and stopping less often. Energy offers the physiological equivalent: a direct, functional read-out of whether the system as a whole is improving, rather than an isolated pathway.

This also addresses a long-standing obstacle. Because energy has been hard to measure precisely, it has struggled for scientific recognition, even as medicine has accumulated biomarkers that can be measured frequently and to several decimal places yet say little about the state of the whole system.

Constraint-Focused P5 Medicine deliberately favors an approximately-right, system-level lead measure over a collection of precisely-measured variables that, individually, may not reflect the health of the system.

Diagnosis, reimagined as learning together. It is worth returning to what the word itself was meant to carry. *Diagnosis* descends from the Greek *diagnōsis* — from *dia-* and *gignōskein*, to know and to learn — and the sense we most want to recover is a collaborative one: not a verdict handed down, but knowledge arrived at *together*. That is precisely what we are proposing.

We are excited by the prospect of bringing the best of human clinical judgment together with rigorous focusing-and-prioritization frameworks such as the Theory of Constraints, with AI able to find pattern across more data than any person can hold (Topol, 2019), and with Digital Twins able to rehearse an intervention before it is tried on a life — each contributing what it does best, none dominating.

The clinician brings context, values and responsibility; the framework brings a disciplined way to move from many findings to the one that matters; AI brings breadth; the Twin brings the ability to test a hypothesis safely. Diagnosis, in this reading, stops being a single expert's pronouncement and becomes what its name always implied — a shared act of learning, undertaken by clinician, patient and machine, and repeated as the constraint moves.

We expect this combination — human judgment and constraint thinking, powered and amplified by AI and Digital Twins — to be used more and more to do two things that were previously beyond any clinician working alone.

The first is to find the *few* locks worth opening: the small number of physiological constraints, hidden among thousands of possible findings like a few needles in a very large haystack, whose relief actually moves healthspan.

The second is to find the *keys* that open them — the specific, best-fit interventions, at the right dose and in the right order, that unlock each constraint. Most of medicine's data tells us about the

hay. What amplified constraint intelligence promises is a faster, more reliable way to find the needles, and the keys that turn.

7. From the Trivial MANY to the Vital FEW and Critical ONE

The general decision problem in healthspan medicine is not a shortage of things that can be measured or improved. It is the opposite: there are so many measurable parts, interdependencies, plausible causes and possible treatments that attention can be dispersed across an almost unlimited field.

How do we distinguish the trivial MANY things we can measure or improve, from the vital FEW that must improve for the system to improve, and from the critical ONE experiment to begin now?

Theory of Constraints contributes two focusing insights.

First, clarity about the system goal separates strong-enough links from weak links. A resource is a constraint only when its capacity, capability or availability is insufficient to meet the demand imposed by the goal. The clearer the goal and required demand, the easier it becomes to identify the constraint set – the vital few weak links – and, where one is more limiting, the critical weakest link.

Second, the many-few-one distinction applies to measurement as well as intervention. We may need many observations to investigate possible causes. We need far fewer measures to know whether the whole system is improving. And we need one measure, or the smallest coherent set, to judge whether the current ONE Experiment is being implemented and producing the expected result.

TRIVIAL MANY	VITAL FEW	CRITICAL ONE
Measurements and improvements that may be real or interesting but do not determine the current decision.	Measures and constraints that reveal whether the system and its weak links are improving.	The experiment and smallest feedback set needed to decide: continue, adjust, stop or reprioritize.

Why the search narrows to one: the fractal nature of Pareto. The Theory of Constraints tells us where to look — the one constraint. Finding the *one problem* beneath it — which cause, or which leverage point, to act on first — rests on a second regularity: the Pareto principle, which is not only real but *fractal*. Suppose a gap has as many as a hundred relatively independent causes. Roughly twenty of them (20%) explain about 80% of the gap; because the pattern repeats at each scale, 20% of those — about four causes — explain around 64% (80% of 80%); and once more leaves roughly a single cause explaining just over half of the whole. This nested structure is why the search can honestly collapse from the one constraint, to the one problem, to the one experiment — instead of fragmenting into many competing explanations.

How do we identify the fewest measurements?

Complex adaptive systems are sensitive to starting conditions. Small differences in one part or interdependency can sometimes produce large changes in system performance, while large changes

elsewhere produce little effect. This makes complete root-cause identification difficult and, in some cases, impossible with the information available.

The difficulty of explaining a system should not be confused with the difficulty of improving it. Some interventions are **cause-robust**: they improve conditions required by several plausible pathways and can therefore be tested before every element of the causal network is known.

Appropriate whole-food nutrition, reduction of excess sugar and highly processed foods, daily movement, restorative sleep and reduction of avoidable stress are familiar examples. They are not universal cures and do not replace diagnosis or treatment of specific disease. Their value is that their effect can be tested directly in the individual.

Three distinct purposes of measurement

Purpose	Question	Typical breadth
1. State and trajectory	Is the person improving, stable or declining relative to the healthspan goal?	A vital few whole-person measures
2. Causes and constraints	Which resources, processes or interdependencies may be unable to meet demand?	Potentially many diagnostic observations
3. Action and learning	Was the intervention implemented, did the constraint respond, and did the person benefit?	The smallest coherent experiment-specific set

Many measurements may be needed to investigate what could be causing a person's current condition. A vital few are needed to manage the system. One measure – or the smallest coherent set – is needed to judge the current experiment.

A system-level measure need not always be technologically sophisticated. One useful signal is the patient's answer to a disciplined question: *Compared with yesterday, last week or last month, do you feel much better, somewhat better, about the same, somewhat worse or much worse?* Subjective experience does not replace objective clinical measures. It is one direct observation of how the whole person is functioning and should be interpreted alongside validated functional, biological and event-based measures.

The ONE Experiment should specify both a **proximal measure**, indicating whether the targeted constraint changed, and a **system-level measure**, indicating whether the person's health, function, resilience or risk trajectory improved. If a local biomarker improves but the person does not, the intervention may have improved a part without relieving the true constraint.

8. The Energy Constraint Hypothesis

If constraints are dynamic and individual, a natural question is whether any candidate recurs often enough to warrant special attention. We propose one, as a strong but explicitly falsifiable hypothesis rather than a universal claim: **energy** — specifically, the body's capacity to produce, distribute and use enough biological energy to meet the demands placed upon it.

This proposal is not free-floating. It answers the question left open by the chaotic topology described in §5: when many links are weak, close together and continuously reordering, the task is to identify the scarce resource under-supplying all of them.

A body in chronic energy deficit is forced into continuous trade-offs between survival, repair and growth, and the systems that lose those trade-offs vary with demand and stressor load – which is precisely why the weakest link appears to move. The instability lies not in the links but in the rationing. On this reading a chaotic constraint topology is a clinical signature of energy deficit, and the therapeutic objective shifts accordingly: not to strengthen any individual system, but to restore a surplus sufficient that the body no longer has to choose.

The rationale is that energy is not one system among many but a substrate on which all others depend. Cellular repair, immune function, protein synthesis, neural activity, hormonal regulation, muscle contraction, cardiovascular function and the maintenance of homeostasis all require a continuous energy supply.

When that supply is insufficient, biology is forced into trade-offs: repair is deferred, immune responses dysregulate, protein turnover slows, muscle is lost, recovery lengthens, cognition suffers and inflammation resolves less readily. Energy behaves, in effect, as the biological currency for which every system competes.

Two age-related trends make this salient.

First, the evidence base on aging indicates that mitochondrial efficiency tends to decline — ATP-generating capacity falls, mitochondrial DNA damage accumulates, oxidative stress rises and mitophagy becomes less efficient — reducing energy *supply*, albeit with wide inter-individual variation.

Second, energy *demand* often rises with age, as accumulated damage requires repair, chronic inflammation taxes the immune system, insulin resistance lowers metabolic efficiency and lost muscle raises the cost of ordinary activity.

The result is a widening gap between supply and demand. Because every cell depends on energy, a sustained deficit can affect metabolism, immunity, endocrine regulation, cognition, physical performance and resilience *simultaneously* — precisely the signature of a system-level constraint. This view aligns with Simpson’s work on mitochondrial medicine and with an expanding literature linking mitochondrial dysfunction to many age-associated disorders (López-Otín et al., 2013, 2023).

The Energy Constraint




Figure 8. The Energy Constraint. Age-related decline in energy supply against rising demand, and the resulting system-wide deficit.

It is essential to distinguish a strong hypothesis from a universal conclusion. We do **not** claim that impaired energy metabolism is the dominant constraint for every individual, nor that all chronic disease reduces to mitochondrial dysfunction. We claim something more modest and more useful: energy is among the **strongest candidate constraints** likely to limit healthspan in many people, especially with age.

Constraint-Focused P5 Medicine therefore does not *assume* energy is dominant — it *asks whether it is*, using diagnostics, wearables, AI, Digital Twins and clinical judgment to test the hypothesis for each person.

For one patient the answer may be impaired energy production; for another, the constraint may lie in immune regulation, sleep, nutrition, psychological stress or musculoskeletal function. The framework's job is not to prove energy is always the answer but to determine the answer for each individual.

For many people in the second half of life, one of the first questions worth asking is simple: *does this person have enough biological energy to meet the demands placed on body and mind?*

Where the answer is no, restoring that balance may be among the highest-leverage opportunities available.

Closing the gap: five leverage points, in order. Once energy is framed as a constraint, health acquires a simple operational definition: supply reliably meets demand, with a small surplus left over. That surplus is the point — it is the body's protective capacity, the slack that absorbs a bad night's sleep, an infection or a hard effort without falling into deficit. Disease begins when demand outruns supply and the buffer is spent. (The word *small* matters: the goal is a slight surplus, not a maximal one; runaway surplus is over-provisioning, its own form of waste.)

If the gap is demand minus supply, there are exactly five ways to close it and build that buffer. The order matters, and it follows two rules. The first four are free and internal; the fifth costs money and

creates a dependency, so it comes last. And within the free levers, the one with the longest lead time is started earliest — which is why creating the conditions for supply sits at number two, even though it is the slowest of the five to pay out.

1. **Stop wasting supply — the efficiency lever.** Capture the energy already produced instead of losing it to friction and oxidative damage. This changes neither how much energy is made nor how much is spent, only how much of what is made actually reaches the work.
2. **Create the best conditions for optimum supply — the generative lever.** Give the body the conditions under which it produces more energy by itself: the right food, enough water, enough of the right movement, enough sleep, and a mind that is not running a standing alarm. This is the only lever that builds capacity rather than redistributing it — supplementation rents energy; conditions build the plant. It sits second, despite being the slowest to pay out, because it has the longest lead time and is the only lever that compounds: it is upstream of the other four, and by Goldratt's own project logic the task with the longest lead time must start first or the whole plan waits on it.
3. **Allocate supply more optimally — the routing lever.** Route what you have to where it is needed, rather than locking it away or leaving the buffer idle. (In this reading, insulin resistance is less a "sugar problem" than an allocation-and-sensing failure: fuel piles up in the wrong places because the routing signal is no longer heard.)
4. **Reduce unnecessary demand — the demand lever.** Shrink the bill rather than the budget: remove the loads the body should not be carrying at all — the standing metabolic tax of ultra-processed food, a dysbiotic gut, stored toxins, or a chronic alarm signal that is still being processed.
5. **Supplement supply from outside — the last-resort lever.** Only if a gap still remains after the first four, add fuel directly (a ketone/BHB drink, MCT oil) to bridge a high-demand moment — used to reach a small surplus and then tapered, not leaned on permanently. Pouring fuel into a system that is still leaking, mis-routing and over-burdened mostly adds to the waste.

This ordering is a sensible default, not a rule to enforce. Because the levers interact — and their payoff depends on the individual — the right sequence is itself something to discover. The digital twin's task is to test these actions and their combinations and report which order produces the largest and most *sustainable* improvement in the whole system. The natural sequence is where to start; the twin tells us where each person should actually go.

The discipline that keeps the generative lever honest is to apply the constraint question one level down: the conditions for production are themselves a chain with a weakest link, so the task is to find the *single* limiting condition — sleep, protein, movement — rather than pursuing all of them at once.

The logic generalizes beyond energy: once *any* dominant constraint is identified, the clinical task becomes closing its specific supply-demand gap, spending the least costly, most controllable lever first and reserving external substrate for last.

If energy is the constraint, improving healthspan reduces to three hard problems, and they map onto the physician's three responsibilities in the room.

The first is **measurement** — how to gauge energy demand against supply across an ordinary day; here we heed Goldratt's caution against being precisely wrong, using subjective report and a handful of proxies until supply and demand can be measured directly.

The second is **decision** — choosing not merely a few reasonable actions but the *one* best lever, at the right dose, to move a person from frequent deficit to reliable surplus.

The third is **inertia** — helping a real person actually act, which we argue (§10) yields not to more willpower but to resolving the underlying conflict. Measurement, decision and inertia on the page; test, prescribe and adhere in the clinic — the same short list in a clinician's coat. The first is a matter for Constraint Intelligence; the second and third are the subject of what follows.

Energy may not be the only such candidate. A second, which we hold more tentatively, is connection: both the biological coupling through which systems signal and regulate one another, and the relational and existential connection, to other people, to purpose and to meaning, that shapes autonomic tone, inflammatory load and the will to act. Where energy is the currency every subsystem spends, connection may be the wiring that determines how well the whole coordinates. We name it as a second falsifiable candidate for a true system-level constraint, to be tested with the same discipline as the first, and note that in a complex adaptive system the deepest limit on a hundred-year healthspan may prove to be some interaction of the two.

9. The Demand Constraint: When the Body Is Not the Limit

Every case so far has assumed the limit lies somewhere inside the body — in one organ system (a simple system), in a few interacting systems (a complex system), or in energy itself, the shared currency whose shortfall starves every subsystem at once (a chaotic system). But there is a fourth case our framework must address: the person whose organ systems are, by every measure, good enough — able to meet daily demand and to last a hundred years or more — and who is declining anyway. For them the binding constraint is not supply of any kind. It is *demand*. The presenting symptom often looks like a connection constraint — withdrawal, isolation, a fading will to act — but the root problem, on inspection, is a lack of meaning and purpose: the absence of feeling wanted and needed. This inverts the usual clinical question.

The clearest way to see it is by analogy to a business, where the same five-constraint structure appears. Ask what limits an organization's ability to reach its goal, and the honest answer is one of five things: too little **demand** for what it offers; too little internal **capacity** to meet the demand it has; too little external **supply** of the inputs it needs; too little **cash** to fund the gap between spending and being paid; or too little **management attention** to direct the whole. The Theory of Constraints tells us that only one of these binds at a time — and that the way to find it is not to guess but to ask them in a deliberate order.

That order begins with the goal and works outward. Is the **goal** clear and worth achieving? Then, is there enough **demand** to reach it? If demand is sufficient, does the business have the internal **capacity** to convert that demand into delivery? If capacity is adequate, is there enough external **supply** of materials, talent and inputs to feed it? And only if all of these hold does the question become **cash**: whether the enterprise can fund itself through the gap. Cash is the quiet chaotic constraint — when it runs short it does not starve one department, it undersupplies every part of the business at once, which is why a cash crisis feels like everything failing simultaneously.

The word *capacity* hides a distinction worth making explicit, because a part can fail to meet demand in four different ways. Each part of a system faces not one demand but two: a **daily demand** — the load placed on it in ordinary operation — and a **durability demand** — the requirement to

keep meeting that load, day after day, for years. To satisfy both, a part needs the **capacity** to handle the volume of demand, the **capability** to do what the demand actually requires, the **availability** to be there when the demand arrives, and the **durability** to keep meeting it over time. These come apart in practice. Picture a bright intern with ample capacity, full availability and youthful durability who nonetheless lacks the capability the task demands; and picture a seasoned employee with deep capability who no longer has the capacity, availability or durability to absorb the day's peaks.

This is why the honest test is not average demand but *peak* demand. A part that meets the average load but buckles at the peaks has not met the demand; it has only deferred the failure. And the deficit compounds: a peak that is not absorbed spills into the next period as backlog, the backlog raises the following peak, and what began as a small, intermittent shortfall becomes a widening, self-reinforcing gap. In a business it is the order that ships late and loses the next order; in a body, as we will see, it is the system that cannot meet a stressor and carries the unrepaired damage forward.

These five constraints map onto a human being with almost unsettling directness. The **goal** is a long, healthy, meaningful life. **Demand** is the reason to live it — the sense of being wanted and needed that pulls a person into each day. **Capacity**, read in its full sense, is the capacity, capability, availability and durability of the organ systems to meet the body's daily and lifelong demands. **Supply** is the external substrate — nutrients, oxygen, sleep, social input — that those systems draw on. And **energy** is the human analogue of cash: the shared currency every subsystem spends, the chaotic constraint that, when it runs short, undersupplies all of them at once. Sections 5 through 8 worked through capacity, supply and energy. This section completes the set by taking the first constraint — demand — seriously.

Consider, then, the person whose capacity, supply and energy are all genuinely sufficient — whose body could carry them well past a hundred — and who is nonetheless fading. Adding more supply does nothing, because supply was never the limit. Optimizing energy does little, because energy was never scarce. Their binding constraint sits at the very front of the chain, in demand: they have run low on reasons to keep going. The symptom presents as connection — loneliness, disengagement, a flattening of the will to act — but connection is the visible surface of a deeper shortfall in meaning and purpose. Treating the symptom (prescribing more socializing) without addressing the constraint (restoring a sense of being wanted and needed) is the medical equivalent of running a marketing campaign for a product no one has a reason to want.

The demand constraint therefore asks a different question of medicine than any organ-focused framework can. Not “what is broken?” but “what is this life for?” — and it guards against a subtle error the constraint lens exists to prevent: pouring resources into supply and energy for a person whose limit lies elsewhere, or, conversely, offering meaning to someone whose true constraint is a failing subsystem that first needs repair. As with every other constraint in this paper, the discipline is the same: identify the one that binds before acting, and let the digital twin — not intuition — reveal which of the five is actually holding the person back.

10. Decision Intelligence: From Knowing to Doing

The preceding sections concern physiology. This one concerns a different and equally decisive constraint. Identifying the right intervention does not ensure it is implemented — and physicians routinely know what a patient should do long before the patient is willing or able to do it.

The evidence supporting better nutrition, activity, sleep, stress management and appropriate pharmacotherapy is now substantial, yet the World Health Organization has estimated that adherence to long-term therapies in developed countries averages roughly 50%. Poor adherence drives avoidable illness, premature mortality and waste. Why do informed people so often fail to act in their own long-term interest?

For decades the dominant assumption was an *information deficit* — people did not change because they did not understand. Education followed. Education remains necessary but is rarely sufficient: most smokers know the risks; most people living with obesity understand the value of nutrition and activity; most patients with hypertension appreciate the value of their medication.

Behavioral science explains why knowledge is not enough — immediate rewards outweigh distant benefits, effort is aversive, habits compete with intentions, emotion overrides analysis, and tolerance for uncertainty, discomfort and cost varies widely.

As the science improves, this behavioral layer becomes *more* important, not less. AI may identify the optimal intervention; a Digital Twin may simulate its effects; Constraint Intelligence may correctly name the limiting physiology — yet none of this guarantees consistent implementation over the months and years that meaningful change requires. AI increasingly answers “*what should be done?*”; a second discipline must answer “*how do we help people consistently do it?*”

This reframes Participation, the fourth P. When Hood proposed P4 Medicine, prediction, prevention and personalization were still emerging; today AI, Digital Twins, diagnostics and wearables are rapidly strengthening the first three Ps. That very progress shifts the limiting factor from knowledge toward implementation — the constraint becomes *behavioral*.

Patients need not only information but confidence, clarity, commitment, support and accountability, and practical methods for resolving the internal conflicts that postpone difficult change.

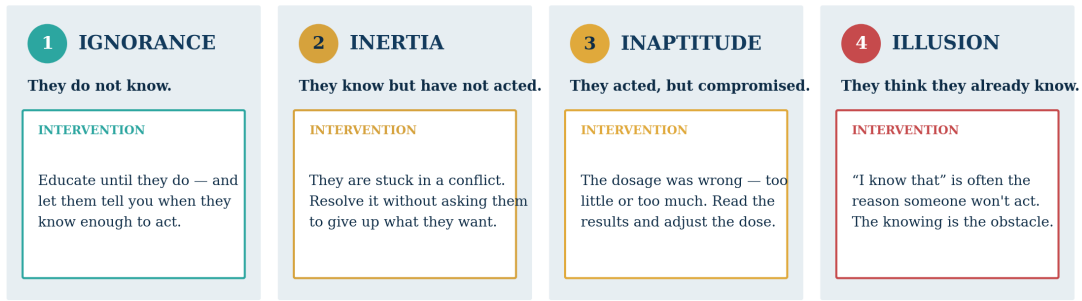
We term the discipline that addresses this **Decision Intelligence**: the systematic improvement of decision *quality* and decision *implementation*. It begins from a different premise than education-as-persuasion — that people generally do the best they can given the decision that seems to make most sense at the time, and that behavior which looks irrational usually reflects an unresolved conflict among competing needs, fears, values or priorities.

Surfacing and resolving those conflicts tends to produce more durable change than information alone. Within Eternity Health this is operationalized through structured mentoring, monitoring and decision-support methods, including the ProConCloud™ method, intended to complement clinical expertise by helping patients understand not only *what* to do but *why* they have struggled to do it before.

Four failures, not one. Framing non-adherence solely as unresolved conflict is itself too narrow. Four distinct failures present identically from the outside and require different interventions; misidentifying which is operating is a common and costly error, because the default response – explaining again – addresses only the first.

The Four I's: Why People Do Not Act

Each failure looks the same from outside and needs a completely different intervention.



Constant mentoring and coaching — from human experts and from the AI coach — works all four.

Figure 9. *The Four I's. Four distinct causes of non-action, each requiring a different intervention.*

Ignorance – they do not know. Education is the correct response, and the endpoint is not the completion of the explanation but the patient’s own statement that they now know enough to act.

Inertia – they know but have not acted. An unresolved conflict is present. It is resolved by finding a route that does not require the patient to sacrifice something they want or accept something they do not.

Inaptitude – they acted but compromised. The dose was wrong: too little to produce an effect, or too much to sustain. The remedy is not motivational but analytical – reading the actual results and adjusting dosage or duration.

Illusion – they believe they already know. “I know that” frequently explains why a patient will not act: the information has been filed rather than used, nothing feels unresolved, and so nothing happens. Of the four this is the one patients are least able to detect in themselves, which is precisely why it requires external observation.

These four are worked continuously, between consultations as well as within them, by the human mentor and by the AI coach. Structured decision-support methods, including the ProConCloud method, address the conflicts characteristic of inertia specifically; they are not a general answer to non-adherence, because three of the four failures are not conflicts at all.

Constraint Intelligence and Decision Intelligence therefore address two different constraints — the *physiological* constraint that limits health, and the *behavioral* constraint that limits implementation. Only when both are relieved can personalized medicine deliver its full value. In this light, participation is not merely asking patients to engage; it is helping them become more capable, more consistent and, ultimately, more successful.

Medicine has long focused on improving diagnosis; healthspan medicine must increasingly focus on improving execution.

The Constraint-Focused P5 Operating System




Figure 10. The Constraint-Focused P5 operating system. Systems biology → diagnostics → wearables → AI → Digital Twins → Constraint Intelligence → Constraint Physiology → CF-P5 Medicine → Decision Intelligence → mentoring & monitoring → healthspan optimization.

11. How Constraint-Focused P5 Medicine Differs

Because the framework deliberately builds on established ideas, it is important to state precisely what is — and is not — new. Constraint-Focused P5 Medicine adds a fifth principle, **Prioritized**, to Hood’s P4 Medicine. The addition does not replace the first four Ps; it *enables* them, supplying the decision layer that determines where predictive, preventive, personalized and participatory effort should be concentrated first.

Table 1. P4 compared with Constraint-Focused P5 Medicine. A dimension-by-dimension comparison of Hood’s four principles with the five-principle model, showing what the addition of Prioritized changes in practice.

Dimension	P4 Medicine (Hood)	Constraint-Focused P5 Medicine
Core question	What is happening across this person’s biology?	Which one process most limits this person’s healthspan <i>now</i> ?
Organizing logic	Measure and understand comprehensively	Measure comprehensively, then prioritize and act on the dominant constraint
Model of the body	Complex biological system	Complex adaptive system governed by shifting constraints
Use of data/AI	Expand what can be observed and predicted	Rank what deserves action; convert observation into a small set of high-value moves
Treatment pattern	Multiple parallel recommendations	Sequential relief of the dominant constraint, then reassessment
Success metric	Coverage and personalization	System-wide improvement per unit of patient effort
Behavior change	Participation (engage the patient)	Decision Intelligence (resolve the behavioral constraint to implementation)

Table 2. P5 Healthcare compared with traditional sickcare. The five principles set against the reactive tradition they build on — offered not as a criticism of a medicine that achieved a great deal without today's testing, AI and digital twins, but as a map of what those tools now make possible.

P5 HEALTHCARE	TRADITIONAL SICKCARE
PREDICTIVE — Anticipates future disease before symptoms develop.	SYMPTOM-DRIVEN — Waits until symptoms become significant enough to investigate and treat.
PREVENTATIVE — Prevents disease by addressing root causes early.	REACTIVE — Responds after disease or damage has already occurred.
PERSONALIZED — Tailors recommendations to the individual's biology, psychology and lifestyle.	STANDARDIZED — Uses standardized guidelines based on population averages.
PARTICIPATORY — Engages patients as active partners in their own health.	DIRECTIVE — Patients are primarily instructed what to do by professionals.
PRIORITIZED — Identifies the ONE intervention most likely to improve healthspan first, then the next.	LIST-BASED — Lists many abnormal findings and fixes without ranking their relative impact.

The framework should not be confused with three adjacent ideas it complements. **Precision medicine** asks why individuals differ in their response to disease and treatment; Constraint Physiology asks a different question — why apparently similar abnormalities produce different consequences in different people — and answers it partly by an abnormality's *position in the constraint hierarchy*.

Root-cause analysis and functional medicine seek an underlying cause of a presenting problem; constraint thinking is distinct in three ways: it is explicitly about *ranking among many simultaneous, interacting factors* rather than tracing one causal chain; it treats the limiting factor as *dynamic*, expecting the priority to change as each constraint is relieved; and it makes prioritization an *operational, testable* procedure (§6) rather than a diagnostic intuition.

Finally, it is not a claim that biology *is* a factory (see §14); it is a claim that one well-validated principle from systems science — that a few constraints govern the performance of the whole — is worth testing rigorously in physiology.

Stated plainly, the novel contribution is not a new measurement, drug or device. It is a **decision discipline**: making the identification of the highest-leverage physiological constraint, and of the behavioral constraint to acting on it, an explicit, prioritized and empirically testable part of clinical practice.

12. Testable Predictions

A framework earns scientific standing by making claims that could be shown to be wrong. We therefore state the core propositions as falsifiable predictions, each with an indication of what would confirm or refute it. These map directly onto the research agenda in §17.

H1 — Dominant constraints exist and are identifiable. In a defined population undergoing comprehensive assessment, individuals will exhibit an identifiable dominant physiological constraint that accounts for a substantial share of their health status. *Refuted if* independent clinicians, given the same data, cannot identify concordant constraints above chance, or if no small subset of factors explains meaningfully more variance in outcomes than the average factor.

H2 — Prioritization can be operationalized. A transparent scoring model (Leverage × Confidence × Feasibility, §6), potentially augmented by AI, will identify dominant constraints with greater inter-rater reliability and predictive validity than unstructured judgment. *Refuted if* the model does not outperform clinician intuition or simple baselines.

H3 — Prioritized care outperforms comprehensive parallel care. In a randomized comparison, care that sequentially relieves the identified dominant constraint will produce greater improvement in validated healthspan outcomes — at equal or lower patient burden — than comprehensive, non-prioritized care. *Refuted if* prioritized care shows no advantage in outcomes or burden. *(This is the framework's central clinical test.)*

H4 — Constraints are dynamic. After a dominant constraint is relieved, a different constraint will more often become limiting, in a detectable and clinically meaningful sequence. *Refuted if* the initial priority remains dominant regardless of intervention, or if changes are random rather than systematic.

H5 — The Energy Constraint Hypothesis. Among aging adults, a measurable imbalance between energy supply and demand will be the dominant constraint in a substantial (though not universal) subset, and relieving it will yield broader system-wide benefit than interventions targeting narrower pathways. *Refuted if* energy status rarely emerges as dominant, or if relieving it produces only pathway-local effects.

H6 — The behavioral constraint is rate-limiting once physiology is prioritized. Once the correct physiological priority is identified, structured Decision Intelligence (mentoring, monitoring, ProConCloud) will improve long-term adherence and outcomes more than information/education alone. *Refuted if* decision-support adds no durable adherence benefit over standard education.

13. Anticipated Objections and Responses

A framework this ambitious should confront its strongest criticisms directly.

“Human biology is not a factory; the Theory of Constraints does not belong in physiology.” We agree that biology is vastly more complex, redundant and adaptive than any engineered system, and we do not import the mechanics of the Theory of Constraints wholesale. Nor do we claim that every patient has a single weakest link. §5 sets out three constraint topologies and a different strategy for each: the single weakest link where one is clearly separated, the few weakest where prediction uncertainty makes them inseparable, and – in the chaotic case – no link-level treatment at all, but the higher-order resource starving all of them. What we adopt is one general principle of complex adaptive systems, that a few constraints disproportionately govern the whole, and we treat its applicability to physiology as an empirical question (H1, H3). The analogy is a source of hypotheses, not of conclusions.

“This is just root-cause analysis (or functional medicine) rebranded.” It overlaps but differs in three respects (§11): it ranks among many simultaneous interacting factors rather than tracing a single cause; it treats the dominant factor as dynamic and expects it to change; and it makes prioritization an explicit, testable procedure rather than a diagnostic intuition.

“The weakest link is usually obvious, so this adds little.” Sometimes it is obvious — and then the framework simply formalizes good practice. Often it is not: abnormalities interact, the loudest

signal is not always the most limiting, and comprehensive assessments routinely generate long undifferentiated problem lists. The contribution is precisely in the non-obvious, multi-factor case, and in making the reasoning explicit and improvable.

“Prioritizing one constraint risks neglecting serious problems.” Prioritization is not exclusion. Constraint-Focused P5 Medicine presupposes comprehensive assessment and standard-of-care management of urgent or dangerous findings; secondary factors are *deferred and monitored*, not ignored, and are revisited at each reassessment (§15, Appendix B).

“Choosing a single constraint could oversimplify and cause harm.” This is a real risk, and the framework is explicitly falsifiable partly to guard against it: H3 tests outcomes *and* burden against comprehensive care, and the approach retains clinical judgment, safety-netting and reassessment. If prioritized care underperforms or causes harm, the framework predicts its own rejection.

“The Energy Constraint Hypothesis privileges one answer.” It is offered as a testable candidate, not a default. The method begins by asking whether energy is limiting for a given person and is designed to identify other constraints when it is not (§7, H5).

“‘Prioritized’ as a fifth P is marketing, not science.” The label matters less than the claim behind it, which is specific and testable: that explicit, individualized prioritization improves outcomes per unit of effort (H2, H3). The framework should be judged on those results, not on the name.

14. Limitations

Several limitations warrant emphasis. First, this is a *conceptual* framework: its central clinical claim (H3) has not yet been tested in controlled trials, and its constructs are working proposals, not validated instruments.

Second, “constraint” is not yet operationally defined at the level of measurement — distinguishing a genuine limiting *constraint* from a correlated *risk factor* is itself an open methodological problem.

Third, the candidate taxonomy (§5) and the prioritization heuristic (§6) require empirical calibration and may be revised or replaced.

Fourth, the Energy Constraint Hypothesis, though supported by converging evidence, remains unproven as a *dominant* constraint at the individual level.

Fifth, AI and Digital Twin components are early-stage and carry known risks — data quality, bias, over-reliance and interpretability — that must be managed.

Finally, prioritization introduces its own risk (premature closure on a single factor) that must be controlled through comprehensive assessment, safety-netting and reassessment. We regard these limitations not as reasons for hesitation but as the specification for a research program.

15. Implications for Clinical Practice

A framework must ultimately improve practice, and should be judged by usefulness rather than elegance. Four implications follow.

The physician's role shifts toward prioritization. With more knowledge available than any clinician can apply, one of the physician's most valuable contributions becomes establishing clinical priorities — helping the patient understand which constraint most limits their health now. This *raises* the value of comprehensive assessment, because identifying the dominant constraint requires understanding the whole system; but assessment is followed by focused action rather than diffuse intervention.

Assessments end in a priority, not a problem list. Traditional evaluations often conclude with many diagnoses and recommendations whose combined burden exceeds what a patient can implement, so little changes despite excellent advice.

Constraint-Focused P5 Medicine asks a single, powerful question after assessment: *"If we could improve only one aspect of this person's physiology over the next three to six months, which would most likely produce the greatest overall benefit?"* Secondary factors are recognized and monitored; once the dominant constraint is meaningfully relieved, the patient is reassessed and the next priority addressed.

Patients are typically more motivated by visible progress on one front than by many simultaneous, individually modest changes.

A common organizing principle for multidisciplinary care. Healthspan care increasingly involves physicians, dietitians, exercise physiologists, psychologists, nurses, coaches and physiotherapists. A shared dominant constraint lets each discipline align around the same primary objective rather than pursuing independent priorities. AI and Digital Twins strengthen this by improving the quality of shared information, while clinical judgment remains responsible for what is appropriate given the patient's history, values and goals.

A different physician-patient relationship. The consultation moves from delivering instructions to building shared understanding — why one intervention takes priority, and how relieving that constraint may improve several aspects of health at once. This strengthens engagement and adherence, and gives concrete form to the participatory principle at the heart of P4 Medicine.

Reports should name the mistake being avoided. Every step of the clinical process can fail in two opposite ways, and most reports never say which one the clinician was working to avoid. Assessment can offer false reassurance or a false alarm; diagnosis can name the wrong cause or miss the right one; treatment can prescribe what will not work or omit what would; a dose can be too much or too little; a follow-up plan can promise too much too soon or set no expectation at all. Because avoiding one error tends to increase the risk of the other, a patient deserves to know not only the conclusion but the trade-off behind it — "we considered telling you not to worry, and we considered alarming you; here is why we landed in the middle." Making that reasoning explicit, and pairing each recommendation with an experiment — what should change, by when, and what would falsify the hypothesis — turns a static report into a feedback loop the clinician can learn from and the patient can trust.

A system-wide win: reducing overwhelm and distrust. The overwhelm created by information abundance is not the patient's alone.

Clinicians and administrators face the same escalation — an expanding menu of tests to order, diagnoses to consider and interventions to choose among, across nutrition, exercise, hydration, sleep,

supplementation, pharmacology, medical procedures and behavioral or cognitive therapy — and the decision does not end at *which* intervention, but extends to *what dose* and *for how long* before improvement should be expected.

A prioritized framework narrows this unmanageable field to a defensible starting point, and in doing so offers a win for every stakeholder in the system:

- **Doctors** — confidence in which tests, diagnosis, treatment, dose and duration to choose, and what to measure to know it is working.
- **Patients** — less overwhelm and more trust in those recommendations, and the ability to both measure and *feel* whether a change is helping, not helping or harming.
- **Insurers** — a more consistent, evidence-based basis for what to approve.
- **Pharmaceutical developers** — a faster, lower-cost path to market without compromising standards.
- **Researchers** — clearer criteria for which questions are most worth pursuing.
- **Governments** — a clearer route to improving population healthspan.

These stakeholders today often operate with considerable mutual distrust. A shared, constraint-based logic — grounded in objective, individualized evidence — can reduce both the overwhelm within each group and the distrust between them.

Digital Twins make this tangible: an insurer could base an approval on simulations run on a client's personalized Digital Twin, comparing the doctor's recommended treatment against no treatment and against alternatives — moving approval from adversarial negotiation toward shared, fact-based evidence.

16. Monitoring the ONE Experiment

Monitoring is not simply the act of checking whether a final health outcome improved. A valid feedback loop must distinguish among three different questions: was the prescription implemented, did the identified constraint respond, and did the whole person benefit?

COMPLIANCE METRICS	LEAD METRICS	LAG METRICS
Was the experiment implemented? Track the actual dose, frequency, duration and quality of implementation.	Is the constraint being relieved? Track proximal changes expected before the final outcome appears.	Did the person improve? Track health, function, wellbeing, resilience and risk outcomes.
Examples: medication doses taken; exercise sessions and intensity; sleep timing; dietary adherence; completion of agreed practices.	Examples: sleep efficiency; strength; VO2 max; glucose control; energy reserve; the specific functional or biological marker targeted.	Examples: symptoms; patient-reported wellbeing; physical function; body composition; clinical events; healthspan and lifespan risk indicators.

Without compliance metrics, lack of improvement is uninterpretable. The treatment may have been inappropriate; the dose may have been too low; the duration may have been too short; or the intervention may not have been implemented fully enough to test it. A clinician may otherwise

conclude that the treatment failed when the more accurate conclusion is that the hypothesis was never adequately tested.

A valid ONE Experiment should therefore define the minimum implementation threshold required for interpretation. If compliance falls below that threshold, the experiment has not necessarily failed; it has not yet been adequately tested.

The monitoring discipline can be summarized in three questions:

1. Was it done?	2. Did the constraint respond?	3. Did the person benefit?
-----------------	--------------------------------	----------------------------

Lag metrics tell us whether the patient improved. **Lead metrics** tell us whether the constraint is being relieved. **Compliance metrics** tell us whether the prescribed experiment was actually performed.

Without all three, treatment failure cannot be distinguished from dosage failure or implementation failure.

This monitoring structure also connects directly to the Four I's.

Low compliance may reflect ignorance, inertia, an implementation or dosage error, illusion, or an external barrier such as cost, pain, cognitive impairment or social conditions. The response should be matched to the actual implementation constraint rather than defaulting to more explanation.

17. A Research Agenda

Constraint-Focused P5 Medicine is the beginning of a research program, not its conclusion; it raises more questions than it answers, which we regard as a strength. The agenda maps onto the hypotheses of §12.

Existence and identifiability (H1). Do individuals exhibit identifiable dominant constraints? How stable are they, and under what conditions do they change? Can independent clinicians identify concordant constraints from comparable data?

Measurement and method (H2). Which combinations of laboratory, wearable, imaging and functional data most reliably identify constraints? Can AI recognize limiting patterns invisible to conventional analysis, and can Digital Twins simulate intervention responses accurately enough to inform decisions? How accurately can Digital Twins, applied with a Theory of Constraints lens and expert oversight, identify the constraint whose relief most improves the whole system, and prioritize interventions to relieve it to the level of good enough?

Comparative effectiveness (H3). Do randomized comparisons show that prioritized, sequential care improves validated healthspan outcomes — at equal or lower burden — versus comprehensive parallel care? This is the decisive clinical test.

Dynamics (H4). Do constraints shift predictably as they are relieved, and can reassessment intervals be optimized?

The Energy Constraint Hypothesis (H5). When does energy imbalance become dominant? Which biomarkers best reflect *functional* energy availability rather than isolated aspects of

mitochondrial biology? Do energy-directed interventions produce system-wide rather than pathway-local benefit?

Behavior and implementation (H6). How do structured mentoring, monitoring and methods such as ProConCloud affect long-term adherence once the physiological priority is set?

Health-economic evaluation. Does prioritization improve the efficiency of resource use and long-term outcomes? Given the global burden of chronic disease, the possibility of better outcomes at lower burden deserves rigorous economic study.

Validated healthspan endpoints (functional, biological-aging, patient-reported and event-based), pre-registration, and comparison against strong non-prioritized controls should be standard across this program.

18. Conclusion

Every generation of medicine inherits extraordinary gifts and extends them. Bernard and Cannon established internal regulation and homeostasis; Osler insisted on caring for the person; Hood reimaged medicine as systems biology and P4, prompted by Larry Page to add participation as its fourth principle; Price demonstrated scientific wellness; Furman illuminated inflammation and immune aging; Goldratt revealed how constraints govern complex systems. The next step, we believe, will again come through integration rather than replacement.

Systems biology, advanced diagnostics, wearables, AI and Digital Twins are transforming our ability to observe health. Their greatest contribution may not be the *quantity* of information they produce but their capacity to help clinicians identify, with growing confidence, the constraint that matters most for each individual — and then to help patients overcome the behavioral constraint to acting on it.

Constraint-Focused P5 Medicine proposes that the future of healthspan medicine depends on making care not only more predictive, preventive, personalized and participatory, but also more **prioritized**.

The fifth P does not replace the first four; it enables them.

We offer this framework not as a conclusion but as an invitation — to clinicians, scientists, engineers, behavioral researchers and patients — to test, challenge and improve it.

Its every major proposition, from Constraint Physiology and Constraint Intelligence to the Energy Constraint Hypothesis and Decision Intelligence, is meant to be investigated and, where necessary, revised or rejected. That is not a weakness of the framework; it is the mechanism by which it, and the field, will advance. The question we hope it makes routine is simple:

What is the primary physiological constraint currently limiting this person's ability to live a longer, healthier life — and how can we most effectively help them overcome it?

If that question yields better research, better conversations and better outcomes for patients, the paper will have served its purpose.

Epilogue: Standing on the Shoulders of Giants

No framework emerges in isolation, and we want to say plainly that we regard our own work as standing on the shoulders of giants. Constraint-Focused P5 Medicine is not a departure from modern medicine but a continuation of its evolution, and we acknowledge with respect the pioneers whose work made ours possible: Bernard and Cannon, Osler, Hood, Page, Price, Furman and Goldratt.

We owe the deepest debt to Dr. Eliyahu (Eli) Goldratt — mentor to one of us — who left us with words at once haunting and inspiring: *“Do not dare stand in my shadow; stand on my shoulders.”* This paper is an attempt to honor that instruction — standing on his shoulders, and on those of the giants named here.

Our debt to Dr. Leroy Hood and Dr. Nathan Price is particularly direct. The fifth P we propose has no meaning without the four that Dr. Hood defined. Prioritized is an addition to P4 Medicine, not an alternative to it, and it exists only because prediction, prevention, personalization and participation had already been articulated and defended. Dr. Price’s work on dense, longitudinal data and, with Dr. Hood, on scientific wellness in *The Age of Scientific Wellness* (2023) showed that wellness itself can be measured, modeled and acted upon, which is the ground our prioritization method stands on. We offer this paper in that spirit, as one more step along a path they opened.

We extend that acknowledgment to Larry Page. The fourth P exists because he insisted that the patient is a participant rather than a recipient, and every prioritized plan we describe depends on a patient who chooses it. A framework about what to do first is worth little if no one does it.

We are grateful, too, to those who opened doors and sharpened our thinking as this work took shape. Blair LaCorte — entrepreneur and Vice Chairman of the Buck Institute for Research on Aging — introduced the authors to the Buck Institute, where Dr. Eric Verdin (MD, PhD), its President and CEO, engaged generously with our energy hypothesis and pointed us to the work of Dr. David Furman — Professor and Director of the Buck Institute’s AI and Bioinformatics Platform and Head of the Stanford 1000 Immunomes Project at the Stanford School of Medicine — on immune aging, and of Dr. Martin Picard on mitochondrial function.

We thank Dr. Donald Shrum, among the first to challenge us to apply the Theory of Constraints to the human body, and Dr. Azra Raza, whose early encouragement expanded this work and who introduced one of us to the ideas of Dr. Daniel Kahneman. And we acknowledge Dr. Herbert Simon, whose insight — that as information grows abundant, attention itself becomes the scarce resource — is part of the foundation for moving from many, to few, to one.

Our two journeys — one beginning in medicine and metabolism, the other in systems science and decision-making — developed independently for years. Only recently did we come to see that they may be exploring different dimensions of one question: *what ultimately limits human health?*

We do not claim to have answered it. We hope, instead, to have framed it well enough that others will stand on our shoulders, improve these ideas, and carry them further than we could have imagined.

Appendix A: Working Definitions

Healthspan. The period of life during which an individual maintains physical, cognitive, emotional and social wellbeing with minimal impairment from chronic disease or disability.

Constraint Physiology. The study of how limiting physiological constraints influence health, resilience, adaptation, aging and disease. Unlike traditional physiology, it focuses on identifying the process currently limiting the overall performance of the biological system; it complements, and does not replace, conventional physiology.

Physiological constraint. The biological process that currently exerts the greatest limiting influence on an individual's capacity to maintain or restore health. Constraints are dynamic and may change as health improves or deteriorates.

Constraint Intelligence. The disciplined process of identifying an individual's dominant physiological constraint by integrating clinical expertise, systems biology, advanced diagnostics, wearables, AI, Digital Twins and longitudinal data. It is a decision-support framework, not a diagnostic device.

Constraint-Focused P5 Medicine. An extension of P4 Medicine that adds a fifth principle — Prioritized — recognizing that sophisticated measurement requires equally sophisticated methods for determining where clinical attention should be directed first.

Decision Intelligence. The discipline concerned with improving the quality of decisions and the consistency of their implementation; within healthspan medicine, it addresses the behavioral constraint that prevents sustained action even when the right action is understood.

Appendix B: A Clinical Workflow

1. Understand the individual.	Clarify the person's goals, history, context, demands, values and the dual goal of survival and growth.
2. Measure comprehensively enough to identify risk and candidate constraints.	Use clinical, functional, biological and patient-reported data to understand state, trajectory and possible weak links. Comprehensive does not mean measuring everything that can be measured.
3. Separate the diagnostic MANY from the monitoring FEW.	Use broad measurement to investigate possible causes, then select the vital few measures that best represent whole-person progress, constraint sufficiency and safety.
4. Identify the constraint set and current weakest link.	Determine which resources cannot meet demand because of insufficient capacity, capability or availability. Where possible, identify which constraint should be addressed first.
5. Design the ONE Experiment.	State the intervention, dose, frequency, duration, minimum compliance threshold, expected lead response, expected lag outcome, review date and stopping rule.
6. Support implementation.	Use Decision Intelligence, mentoring and monitoring to address ignorance, inertia, implementation error, illusion and external barriers.
7. Monitor three metric types.	Track compliance metrics (was it done?), lead metrics (did the constraint respond?) and lag metrics (did the person benefit?).
8. Learn and reprioritize.	Interpret the three metrics together. Continue, adjust, stop or reprioritize. When a constraint becomes strong enough, identify the next limiting constraint or constraint set.

**The central learning loop is: MEASURE → PRIORITIZE → ACT → MONITOR
COMPLIANCE, LEAD AND LAG → LEARN → REPRIORITIZE.**

Disclosure

Author affiliations. The authors are affiliated with Eternity Health AI, a commercial entity developing clinical applications of the framework described in this paper.

Author contributions. The framework is the product of two complementary bodies of work. Dr. Graham Simpson contributed the clinical foundation: the *Era III Medicine* framework, the four primary hallmarks of aging, and the recognition of mitochondrial function as the constraint on the body's capacity to generate enough energy to meet demand — together with the clinical content throughout the paper. Dr. Alan Barnard contributed the systems and decision-science foundation: the naming of *Constraint-Focused P5 Medicine*; the three types of link in a chain (strong-enough, weak and weakest); the three system topologies (simple, complex and chaotic); all insights drawn from the Theory of Constraints; the *ONE Thing Focusing Cycle*; the Four I's; the extension of VUCA to VUCCA, adding constraints as a fifth condition; the framing of energy demand outrunning supply; the three classes of monitoring measure (lead, lag and compliance metrics); and the architecture of a personalized Digital Twin built from a patient's clinical-test data.

Competing interests. Constraint-Focused P5 Medicine and its associated methods form the conceptual basis of Eternity Health AI's clinical program, and the authors therefore have a financial interest in its adoption. *ProConCloud* and *Harmony Decision Maker* (which guides users through the 5-Step ProConCloud method) are trademarks of Dr. Alan Barnard. The remaining concepts introduced here — including Constraint Physiology, Constraint Intelligence, Decision Intelligence and Constraint-Focused P5 Medicine — are offered as open scientific concepts intended for free use, discussion and testing by the wider research and clinical community.

Funding. No funding was received from any organization to complete the research for this paper.

References

- Barnard, A. (2026). *From Many to Few to One: The ONE Thing Focusing Cycle to Do More, Faster, Better, Cheaper and Simpler*.
- Bennis, W., & Nanus, B. (1985). *Leaders: The Strategies for Taking Charge*. New York: Harper & Row.
- Bernard, C. (1865). *An Introduction to the Study of Experimental Medicine*.
- Cannon, W. B. (1932). *The Wisdom of the Body*. New York: W. W. Norton.
- Furman, D., Campisi, J., Verdin, E., et al. (2019). Chronic inflammation in the etiology of disease across the life span. *Nature Medicine*, 25(12), 1822–1832.
- Goldratt, E. M. (1990). *What Is This Thing Called Theory of Constraints?* Great Barrington, MA: North River Press.
- Goldratt, E. M., & Cox, J. (1984). *The Goal: A Process of Ongoing Improvement*. Great Barrington, MA: North River Press.
- Hood, L., Heath, J. R., Phelps, M. E., & Lin, B. (2004). Systems biology and new technologies enable predictive and preventative medicine. *Science*, 306(5696), 640–643.
- Hood, L., & Friend, S. H. (2011). Predictive, personalized, preventive, participatory (P4) cancer medicine. *Nature Reviews Clinical Oncology*, 8(3), 184–187.

- Hood, L., & Flores, M. (2012). A personal view on systems medicine and the emergence of proactive P4 medicine: predictive, preventive, personalized and participatory. *New Biotechnology*, 29(6), 613–624.
- Hood, L., & Price, N. D. (2023). *The Age of Scientific Wellness: Why the Future of Medicine Is Personalized, Predictive, Data-Rich, and in Your Hands*. Cambridge, MA: Belknap Press of Harvard University Press.
- Hood, L. (n.d.). Quoted on the origin of the fourth P (participatory) in conversation with Larry Page. PubMed Central, PMC2957728.
- Kitano, H. (2002). Systems biology: a brief overview. *Science*, 295(5560), 1662–1664.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2013). The hallmarks of aging. *Cell*, 153(6), 1194–1217.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2023). Hallmarks of aging: an expanding universe. *Cell*, 186(2), 243–278.
- Price, N. D., Magis, A. T., Earls, J. C., et al. (2017). A wellness study of 108 individuals using personal, dense, dynamic data clouds. *Nature Biotechnology*, 35(8), 747–756.
- Simpson, G. (2025). *Era III Medicine: One Mind and the Ever-Present Origin*.
- Topol, E. J. (2019). *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. New York: Basic Books.
- Weston, A. D., & Hood, L. (2004). Systems biology, proteomics, and the future of health care: toward predictive, preventative, and personalized medicine. *Journal of Proteome Research*, 3(2), 179–196.
- World Health Organization. (2003). *Adherence to Long-Term Therapies: Evidence for Action*. Geneva: WHO.