

From Static Evaluation to Adaptive Decision-Making: Is a Real Options Framework Suitable for Healthcare?

Michael Ghilissen

Department of Public Health Sciences, Faculty of Medicine, University of Liège

Michael.ghilissen@uliege.be

Introduction

Healthcare organizations operate in environments shaped by clinical, regulatory, operational, epidemiological, and financial uncertainty. Traditional decision-support methods — cost-effectiveness analysis (CEA), cost-benefit analysis (CBA), net present value (NPV), and budget impact analysis — remain indispensable for evaluating efficiency and economic value, but they largely assume stable information and predetermined decision pathways. That assumption undervalues flexibility when evidence evolves, regulations shift, clinical outcomes remain uncertain, and disease progresses irreversibly.

This article proposes treating healthcare choices as sequential commitments rather than one-time decisions, using Real Options Analysis (ROA) as a starting point. Real options thinking highlights the value of waiting, learning, staging investments, and adjusting strategy as new information arrives. Applied to healthcare, though, the framework needs to be reframed: not as a financial options market, but as a dynamic learning system in which staged investment, evidence generation, and adaptive commitment form the basis for governance.

Why the Financial Options Analogy Breaks Down

ROA borrows its structure from financial options theory, treating decisions as calls and puts. As a teaching device, the analogy is useful; as an operational model, it oversimplifies a reality that is multidimensional, evolving, and often irreversible. The typical adaptation maps clinical or project value to a stock price, investment cost to a strike price, healthcare uncertainty to volatility, and a clinical time window to an expiration date.

That mapping holds conceptually but collapses mathematically. Healthcare lacks tradable assets, observable market prices, replication portfolios, and risk-neutral valuation — the mechanisms that underpin options pricing formulas in finance. Healthcare is better understood as a sequential decision system than as an options market, and the argument here is that treating it as the latter yields calculations that look rigorous but are not sound.

Decomposing Uncertainty Into Distinct Domains

A recurring error in real options applications is compressing all uncertainty into a single volatility parameter. Healthcare uncertainty is not a single factor; it comprises several largely independent risks bundled together, and a workable framework needs to separate them:

- Clinical volatility — efficacy, safety, durability, and patient adherence (e.g., Phase II oncology results with wide confidence intervals).

- Regulatory volatility — approval and labeling decisions and post-marketing surveillance; a therapy can post strong trial results and still fail to reach patients.
- Market volatility — reimbursement, competitor entry, provider adoption, and pricing pressure.
- Epidemiological volatility — disease prevalence, viral mutation, and population-level trends, as pandemic planning illustrates.
- Operational volatility — staffing, supply chains, infrastructure readiness, and technology implementation, as in AI deployment inside a health system.

Financial models generally assume symmetric uncertainty, but healthcare outcomes are often skewed. Gene therapy illustrates the range: outcomes can range from curative breakthroughs to moderate benefit to complete failure to severe adverse events. A single variance figure cannot capture that spread. Decision-makers instead need full outcome distributions, tail risks, and catastrophic scenarios — tools such as Monte Carlo simulation, decision trees, and scenario analysis — rather than a single volatility number.

Learning Under Uncertainty, Not Just Waiting

Healthcare decisions are not simply made under uncertainty — they are made as that uncertainty is actively reduced. Drug development illustrates this: each trial phase updates the probability of success based on new evidence, so decision-makers continuously revise a Bayesian estimate. Under this framing, the value of learning often exceeds the value of waiting on its own.

This is also where conventional real options thinking can mislead. Financial options generally benefit from waiting because the underlying asset's volatility tends to work in the holder's favor over time. Healthcare often does not share that property: in stroke, sepsis, or cancer progression, delay can destroy the opportunity a decision-maker is trying to protect. The relevant comparison is the value of additional information against the cost of clinical deterioration — the option to defer is worth exercising only when the former outweighs the latter.

Financial options also expire on a fixed date, whereas healthcare opportunities more often fade gradually. In cancer, the probability that treatment remains effective tends to decline over time rather than falling off a cliff at a set expiration. Hazard functions, which directly reflect disease progression, mortality risk, and biological decline, capture this better than a fixed expiration date.

A Richer Set of Strategic Options

The language of calls and puts is useful for teaching but too limited for operational use. In practice, healthcare decision-makers work with a broader set of choices: the option to defer (wait for stronger evidence before adopting a new device), to stage (commit incrementally, as in sequential trial phases or pilot programs), to expand (scale a successful pilot system-wide), to contract (pull back investment when results disappoint), to switch (move to a different protocol or technology), and to abandon (discontinue a failing drug candidate after Phase II).

Complementing, Not Replacing, Existing Health Economics Tools

Real options analysis is proposed here as a complement to, not a substitute for, QALYs, cost-effectiveness analysis, and health technology assessment — each of which addresses a different question. QALYs measure the health value generated; CEA assesses whether an intervention is efficient; real options assess what flexibility itself is worth. On this view, total healthcare value is the sum of expected health benefit, flexibility value, and learning value.

Where the Framework Applies

The framework fits most naturally where staged investment, high uncertainty, and clear go/no-go milestones already exist — pharmaceutical R&D is the clearest example. It also applies to hospital strategy for AI diagnostics, robotic surgery, and digital health platforms; to public health decisions such as pandemic preparedness, vaccine platform investment, and strategic stockpiles; and to precision medicine, including biomarker-driven treatment pathways and adaptive treatment strategies.

A Caveat: Patients Are Not Financial Assets

None of this should obscure a basic limitation of the analogy: patients are not financial assets. Healthcare goals include survival, quality of life, equity, access, and fairness. Any real options analysis should support clinical judgment and ethical decision-making rather than override them.

Applying the Framework in Hospitals

In practice, the framework is less a valuation exercise than a governance and decision-architecture problem. Rather than asking "what is the option worth," hospitals are better served by asking "how do we structure decisions so we can learn before committing more capital?" That reframing — treating a major initiative as a sequence of options rather than a single investment — makes the approach operational and can be broken down into seven steps.

1. Identify the underlying opportunity before spending anything.

AI radiology is about improved productivity and diagnostic accuracy; robotic surgery is about better outcomes and increased referrals; telemedicine is about lower costs and greater access; a new cancer center is about added capacity and market growth; and predictive analytics is about reduced readmissions. This replaces the "stock price" analogy: the hospital is valuing a future opportunity, not an asset.

2. Map uncertainty instead of collapsing it into one variable.

For AI-assisted radiology, that means asking clinical questions (does accuracy improve, do miss rates fall), operational questions (will physicians use it, does it help or hinder workflow), financial questions (does it affect reimbursement), regulatory questions (what rules are coming), and market questions (will competitors adopt similar systems). That uncertainty map becomes the hospital's volatility profile.

3. Define learning milestones.

The question is what evidence would change the decision. A six-month AI pilot, for example, might track accuracy, reading time, physician acceptance, and cost savings. If those metrics improve, the hospital proceeds; if they don't, it stops. The value comes from learning before scaling, not from scaling and hoping.

4. Build explicit decision gates with predefined investment levels.

A small pilot might carry a €250,000 investment with outcomes of expansion, modification, or termination; a department-wide deployment might carry a €2 million investment with outcomes of hospital-wide scaling, holding steady, or stopping; a system-wide rollout might carry a €10 million investment with outcomes of full implementation, partial implementation, or exit. At each gate, the hospital is buying information before committing more capital.

5. Measure the cost of waiting.

This is where real options discussions often go wrong, because "learning" can become a cover for delay. A sepsis detection platform illustrates the trade-off: waiting buys more evidence and reduces uncertainty, but it also risks avoidable mortality, longer stays, and higher costs. The value of waiting equals the information gained minus the value lost to delay, and the case for moving forward strengthens once delay destroys more value than it creates.

6. Use hazard rates for clinical initiatives.

Clinical opportunities are often perishable rather than fixed to a date. Delaying a cancer screening expansion by a year leads to more late-stage diagnoses, higher treatment costs, and worse outcomes — the opportunity erodes gradually rather than vanishing at a deadline. Estimating the probability of losing the opportunity over time is a better basis for boards than an arbitrary planning horizon.

7. Set expansion and abandonment rules before launch.

Vague exit criteria are a common reason hospital initiatives stall. A program might expand if readmissions fall by more than 10%, ROI exceeds the hurdle rate, and clinician adoption exceeds 70%; it might be abandoned if there is no clinical improvement, cost overruns exceed a set threshold, or adoption falls short. Agreeing to these rules before the pilot starts is intended to keep politics from overriding evidence later.

The contrast between the approaches is well illustrated by three examples. The conventional path for an AI diagnostic tool runs from business case to full purchase to implementation to hoping for the best; a staged path instead runs from a single-department pilot to evaluating outcomes, to expanding to three departments, to evaluating again, to hospital-wide adoption — preserving flexibility at each stage. A surgical robot can follow the same logic: instead of an outright capital purchase, a hospital leases the equipment first, tracks case volume, surgeon adoption, and outcomes, then expands training, and only then buys permanently — the lease functions like an option premium, and the purchase is the exercise decision. A heart-failure readmission program can be piloted with 500 patients, reviewed on readmission rates, mortality, and cost per avoided admission after twelve months, and only then expanded to 5,000 patients, avoiding a multimillion-dollar commitment before the outcomes are known.

Boards typically track NPV, ROI, and budget impact, and these measures remain useful. A real-options dashboard would add flexibility metrics (cost of expansion, cost of contraction, and exit costs), learning metrics (uncertainty reduction, evidence quality, and confidence intervals), clinical metrics (QALYs, mortality, readmissions, and patient satisfaction), and risk metrics (regulatory, technology, and adoption risk).

Conclusion

The central claim of this article is that healthcare uncertainty should not be modeled with a single volatility parameter borrowed from finance. On this view, the value of a real options approach lies not in treating healthcare as an options market but in recognizing it as a process of sequential decision-making, evidence accumulation, and adaptive commitment under uncertainty. A healthcare-specific real options framework would need to move beyond call-and-put analogies and rest on four elements: multidimensional uncertainty, Bayesian learning, irreversibility and disease progression, and strategic flexibility.

The underlying shift proposed here is simple to state and harder to execute; hospitals would stop asking whether to invest and start asking what the smallest investment is that lets them learn whether to invest more. Whether that shift delivers the claimed governance benefits depends on execution — on whether decision gates, hazard-rate estimates, and expansion/abandonment rules are adhered to when pilots produce inconvenient results, rather than renegotiated under institutional pressure.