Worldwide Clinical Trials

If Your Clinical Program Fails, Will You Know Why?

Uncertainty debt and the hidden costs of under-specified trial design

Panteli Theocharous, Ph.D., FRCPath, Chief Therapeutics and Clinical Strategy Officer, President of Europe

Worldwide Clinical Trials | October 2026

Table of contents

Executive summary **3**

Introduction **3**

Uncertainty in therapeutic programs is a debt to be paid **4**

- What makes uncertainty behave like debt? 4

Uncertainty debt compounds and destroys optionality in clinical programs **5**

- The harsh economics of ambiguous negatives 5
- De-risking is about where and when to allocate capital 5
- Irreversibility sets the price of uncertainty 6
- Assurance, not power 6

Where uncertainty debt accumulates across development and modality **7**

- Dose and exposure 8
- Feasibility and protocol conduct 8
- Endpoint and estimand 8
- Uncertainty changes shape by modality 9
- Designing backwards from the label and the price 10

Identifying and retiring uncertainty debt before the next capital gate **10**

- Optimism in clinical studies is structural 10
- A framework for scoring uncertainty 10

Conclusion **12**

About the author **12**

References **13**

Executive summary

When a clinical program fails, the post-mortem usually ascribes the failure to a problem of execution. However, it is more likely that the failure was authored years earlier in design decisions that left decisive questions unanswered. The accumulated cost of those unanswered questions saddles clinical programs with an uncertainty debt. Uncertainty debt is distinct from irreducible biological risk, which must be understood and priced. Uncertainty created by avoidable development and design choices is what ends up as debt, and opportunities should be taken early in the life of a clinical program to engineer it out. Biology determines whether a therapy works. Development design determines whether we understand why a therapy does or does not work.

This paper addresses five crucial questions. Why is failure in clinical trials frequently misdiagnosed? What is uncertainty debt? How does it compound and destroy optionality for clinical programs? Where does it accumulate across therapeutic development and modality? And, how should leadership identify and retire uncertainty before committing further capital to a clinical program? Herein, we explore five ideas that help inform the answers to these questions:

- Irreducible biological risk must be priced, while avoidable uncertainty should be engineered out.
- An ambiguous negative can destroy more capital value for a clinical program than a clean negative, because it prevents decisive action.
- Across the lifetime of a clinical program, the cost of uncertainty rises as optionality falls and decisions become less reversible.
- Uncertainty often carries different consequences across therapeutic modalities.
- There are practical diagnostic questions leadership teams can address before committing further capital to a clinical program.

Introduction: Failure in clinical trial execution is often misdiagnosed

When a therapeutic development program fails, the post-mortem analysis often blames it on a problem of execution. The clinical research organization might be blamed for underperforming. Trial sites might be faulted for poor enrollment. And, of course, the data were messy and the signals couldn't be properly uncovered. Protocol deviations, too, are a popular scapegoat.

Diagnosing failures in study execution is comfortable for two reasons. First, these types of failures are visible and easily picked out, and second, they can be attributed to personnel or the organization. In contrast, a decision-based failure has neither of these attributes. It's invisible in the moment, and by the time it surfaces, the committee responsible for the decision has moved on to other projects, with its reasoning preserved only in documents that will likely never see the light of day again.

As such, what typically happens is that we fix the visible "failure." We change vendors, or add additional monitoring, maybe tighten oversight, or write a new SOP to protect against further deviation. And then, the next program fails the same way.

From 2003 to 2011, the likelihood of drug approval from Phase I was 10.4%.[1] More than a decade later (2014-2023) that number had fallen to 6.7%.[2] The default explanation for that decline is that the biology of emerging therapeutics has become more complicated, and so have the associated clinical studies.

Indeed, an analysis of Phase 2 and Phase 3 failures from 2013 to 2015 attributed 52% to lack of efficacy.[3] This notion of efficacy being the leading cause of failure has hardened into a fact about the industry, but it deserves a second look.

A recent analysis of program terminations (versus clinical study failures) examines the point from another direction. In an analysis of a proprietary industry dataset covering more than three thousand terminated industry Phase 2 and Phase 3 trials between 2013 and 2023, efficacy accounted for only 25% of program terminations, while strategic and business decisions were the leading driver of 34% of program terminations.[4] While it's true that early terminations are not the same population as programs that fail after reading out, the combination of these numbers still presents an uncomfortable truth: an increasing share of what we often call clinical study failure may, in fact, be a program decision made under uncertainty that was never resolved before onset of the clinical trial.

Phase 2/3 trial termination causes:

25% Drug efficacy

34% Strategic & business decisions

As a result, failure surfaces at the readout, when in fact it was authored years earlier in a set of decisions that were reasonably made on the information available at the time, but with a degree of uncertainty that makes failure likely. This isn't an argument that good design rescues bad drugs. It does not, and nor should anyone make that claim. Programs die of clean biological failure all the time, and when they do, the system is working exactly as intended.

However, while biology determines whether a program fails, it's uncertainty that determines what failure costs. A clean failure that returns capital back to the therapeutic pipeline is survivable. A failure resulting from uncertainty in program decision making that consumes capital without resolving anything is frequently much harder for programs to get past. As that uncertainty debt grows, the barrier to the program's success grows with it.

Uncertainty in therapeutic programs is a debt to be paid

Every under-specified clinical study design decision is debt taken against a successful future readout. A missing target-engagement biomarker, carrying forward only a single dose into the next phase, or an unspecified estimand are all examples of uncertainty debt that accrues interest and comes due at the most expensive moment in the study. This liability doesn't appear on the balance sheet, so nobody prices it. As a result, it accumulates without resistance. Decision-relevant uncertainty becomes progressively more expensive as development advances and available options narrow.

What makes uncertainty behave like debt?

Debt is uncertainty we choose not to resolve when we could have. In a clinical study, perhaps the assay existed, the cohort was available, or the extra arm was affordable, but the uncertainty is deferred for reasons that were probably acceptable at the time. But, uncertainty debt has to be repaid. The only question is, "When?" With biological unknowns, we accept them as irreducible with our current understanding and tools, quantify the cost, and make sure everyone funding the program knows they are there. A clinical program that cannot tell these apart may be agonizing over uncertainty it could have dealt with for the cost of a substudy.

What would it cost to resolve uncertainty in clinical study questions sooner than later? For example, an engagement biomarker substudy, an additional dose arm, a natural-history cohort, a potency assay qualified against clinical outcome; these are mechanisms to answer clinical study uncertainty of different kinds. They are almost always seven-figure cost additions, but they exist in clinical programs that will eventually spend nine figures or more.

Every study decision made on uncertain footing compounds. As each milestone passes, a program is not carrying one unresolved question, but instead an entire study structure built on them. At the readout, the uncertainty debt comes due all at once. A question worth four million dollars to answer in Phase 2 might be repaid with a trial at Phase 3 prices, or not at all as the failed asset is written off with the question and uncertainty still open.

Uncertainty debt compounds and destroys optionality in clinical programs

The harsh economics of ambiguous negatives

A clean negative in a clinical trial has an upside. It kills the asset, frees capital and the team, returns management attention to the rest of the portfolio, and preserves whatever was learned from the failure. The cost is characterized and known on the day the database locks. Nobody enjoys the failed asset, but at least it's a cleanly completed program.

An ambiguous negative is far more damaging, because it completes nothing. It does not kill the asset. It funds yet another study. A claim that cannot be cleanly dismissed can cost two more years and a nine-figure trial to adjudicate, while the program asset remains too viable to kill, yet too unclear to advance. I've seen these "undead" assets consume capital, management attention, and credibility as the program is forced to ask investors to fund the resolution of a question rather than the pursuit of a result, two very different propositions to have to defend. The ambiguous negative is why bad news arrives late and with great expense. The information that could have stopped the program cheaply, instead, arrives years later attached to a request for more money.

In my experience, programs survive a clean negative more readily than an ambiguous one. The former permits a decision. The latter demands another investment. For a large organization, unresolved uncertainty may be absorbed within a portfolio. For a biotech with finite capital and runway, it can become existential.

De-risking is about where and when to allocate capital

De-risking is not synonymous with caution. It is the disciplined allocation of capital to resolve decisive questions while the answers can still shift and enhance the program.

In clinical programs, the value of information declines as cumulative spending rises. Answers to questions of uncertainty are worth what they can change and improve, and by late Phase 3, there are very few decisions left to be changed, as is shown in Figure 1. An answer to a source of uncertainty that could have positively redirected the program now, instead, merely explains the impact of that uncertainty two to four years later.

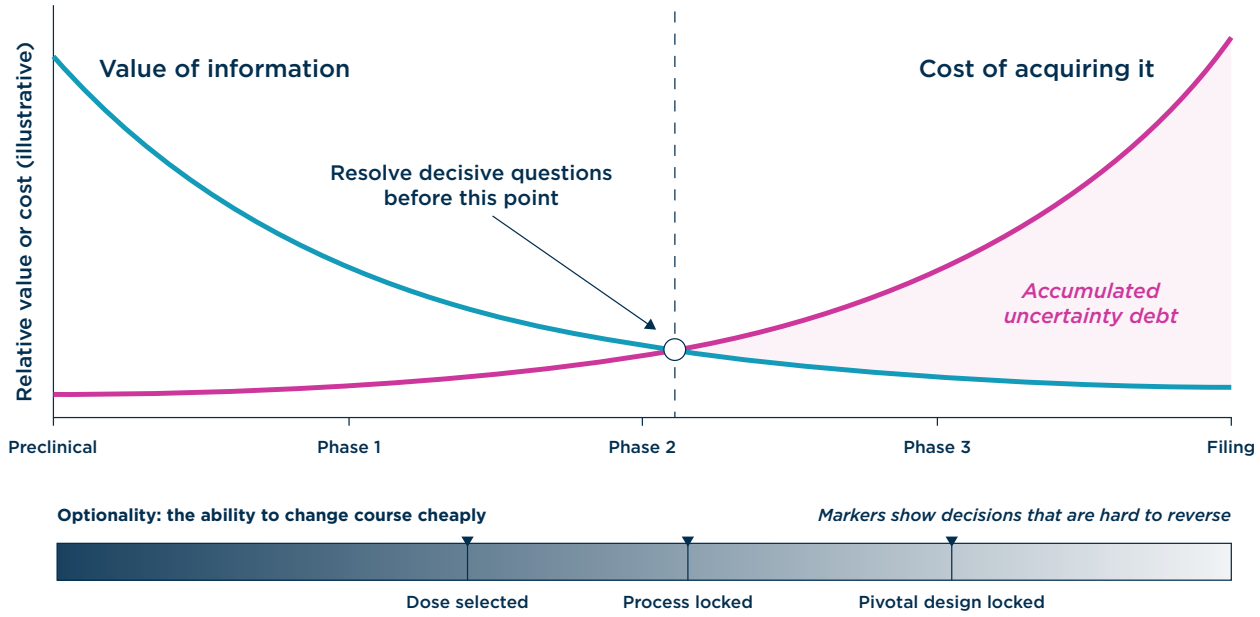


Figure 1. *The cost of uncertainty rises as optionality falls. The value of resolving a question declines as a program advances, while the cost of answering it climbs. Questions left open past the crossover become uncertainty debt. Curves are illustrative and not to scale.*

The cost of acquiring that information also climbs with every phase, for obvious reasons, as the trials get larger, more sites are needed, possible need for longer follow ups are needed, etc.

De-risking a clinical program is about getting ahead of uncertainty before the point where the two curves, representing information value and study cost, cross. An experiment has different value at different times, and the question is never just whether to run a study, but also when. For every proposed study, substudy, or analysis, name what will change if the answer is "yes," and what will change if the answer is "no." If the two are the same action, the question is not decisive and should be deferred, regardless of how scientifically interesting it is.

Irreversibility sets the price of uncertainty

What de-risking affords a clinical program is optionality. For example, futility interim analyses, sample-size re-estimation, adaptive dose-finding, master protocols sharing a control arm, or pre-set kill criteria are elements typically discussed as efficiency measures. Yet, in practice, every one of them allows a program to change direction cheaply. However, that optionality decays as the program becomes more committed to the endpoint, and by mid-Phase 3 much of the potential optionality available to a clinical program has been lost.

As such, the cost of uncertainty is determined by both its magnitude and the reversibility of the decision made under it. In my experience, large reversible commitment is often a safer position than a small irreversible one, and reversibility belongs at the center of how study actions are priced. This principle carries particular weight in advanced therapeutics. For example, in cell and gene therapy development, dose, manufacturing, potency, and comparability decisions may be difficult to revisit, which makes unresolved uncertainty materially more consequential.

The same study protocol can be 90% powered and carry an assurance near 50%, and the assurance number, built on a credible prior, is the more realistic guide to whether the trial will succeed.

That said, I think it's important to state that speed can and does have real value. In high unmet need, it is a moral imperative. There are programs where the correct decision is to accept ambiguity and move forward. Crucially, there is a distinction that helps to resolve this dilemma. Accelerating past a question that is not decisive is efficiency, and clinical studies should do more of it. But, accelerating past a question that can have decisive power does not reduce cost, it only defers it. And deferred uncertainty debt compounds.

Assurance, not power

Here is a place where standard practice in clinical programs can actively mislead those making capital decisions. Power is a conditional probability, assuming a fixed and known true effect size.[5] It states that if the true treatment effect is exactly the value you assumed, you will achieve statistical significance some percentage of the time. It does not, however, say whether the true effect is that value. Yet "the trial is 90% powered" is reported to boards and investment committees as though it were the probability of success. In fact, it is not.

Assurance, sometimes called probability of success or Bayesian predictive power, addresses the category error. Rather than conditioning on a single assumed effect, it works from a specified and honest prior distribution over the true effect, computes power at every point in it, and averages. The output is no longer conditional on a single assumed effect. Instead, it estimates the probability the trial reads out positive, given that prior. Assurance is therefore only as credible as the prior distribution and the assumptions behind it.

The same study protocol can be 90% powered and carry an assurance near 50%, and the assurance number, built on a credible prior, is the more realistic guide to whether the trial will succeed. Power describes performance under an assumed effect; assurance estimates the probability of success across a credible distribution of possible effects. Additionally, as sample size rises, power climbs toward 100%, however, assurance does not. It plateaus at the probability, under the prior, that the drug works. No trial size, no matter how large, can rescue a weak hypothesis, and unlike a power calculation, an assurance calculation makes that more plainly visible to non-statisticians.

Where uncertainty debt accumulates across development and modality

The debt that uncertainty accrues for a clinical study and inflicts on the capital resources of the program is often hard to discern, because where it is inflicted and where it is felt are often greatly separated from each other. For uncertainty to be effectively removed from the study, the decision to engineer it away must be temporally correct. Otherwise, actions taken to remedy the effects that are felt later in the study are reactive, at best, and do not protect capital. Table 1 shows the major sources of uncertainty debt in clinical studies, when they are introduced into a study versus when their effects are felt, and mechanisms to engineer the debt away at the proper time. Three of these sources of uncertainty deserve a deeper understanding.

Table 1: Sources of uncertainty debt in clinical studies and mechanisms to avoid them

Domain	When it's introduced » When the effect is felt	What a negative cannot tell you	How uncertainty is avoided
Target and translation	Candidate selection » Phase 2b/3	Wrong target or missed target	Human genetic and human-tissue causality evidence; an engagement readout fielded before efficacy
Product, CMC, and comparability	Process development » pivotal readout, filing, or inspection	Whether the product, rather than the target or dose, fell short; which data apply to the commercial product	Potency assay qualified against clinical activity; product characterization and manufacturing consistency established; process locked before the pivotal; comparability strategy defined in advance
Dose, exposure, and pharmacology	Phase 1 escalation » Phase 3	No effect or wrong exposure	Exposure-response characterized for efficacy and toxicity; more than one dose taken forward into the pivotal
Endpoint, estimand, and statistical inference	Protocol design and sample sizing » readout or review	What was actually estimated; signals that can be neither trusted nor dismissed	Estimand framework applied at design; endpoint position documented, not assumed; assurance from discounted priors; pre-specified testing hierarchy
Population and enrichment	Eligibility criteria » readout	No effect or diluted effect	Prospective enrichment; criteria tested against real-world data for what each costs in patient pool attrition
Feasibility and protocol conduct	Site selection and protocol drafting » 12 to 18 months later, amendments, and dropout	Whether rescue-site heterogeneity contaminated the primary; whether deviations confounded the per-protocol picture	Performance-based selection; real-world eligibility funnels; rescue triggers defined in advance; the CSR-or-label test on every assessment, with a named owner for total burden
Safety	Nonclinical package and Phase 1 » late-phase aggregate review or label negotiation	Whether an adverse-event signal is drug related or background	Mechanism-based hazards enumerated; background rates established; program-level aggregate monitoring pre-specified
Regulatory alignment	Whether and when to seek advice » review	Whether the position was agreed or inferred	Documented positions from pre-IND, end-of-Phase-2, and Type C meetings
Value and access	Phase 3 design » health technology assessment and launch	Whether the product lacks value or the trial failed to demonstrate it	Target value profile defined; comparator strategy aligned for regulators and payers; patient-relevant and health-economic evidence collected prospectively

Dose and exposure

For many programs, this is a costly form of uncertainty debt, because it affects everything downstream of it. The root cause is this: maximum-tolerated-dose logic was developed for cytotoxic agents, where more drug meant more cells killed and the only limit was what the patient could survive. Applied to targeted, immune, and chronic-use agents, it is often the wrong model, but for two decades, many clinical studies applied it anyway, carrying a single dose into registration on the strength of a 3+3 dose escalation scheme in Phase 1.

However, today's regulatory posture has caught up. The U.S. Food and Drug Administration's (FDA's) guidance on dose optimization in oncology, finalized in August 2024 under the Project Optimus initiative, recommends a randomized, parallel comparison of multiple dosages before or alongside the registrational efficacy work.[6] The guidance explicitly says such a trial need not be powered to demonstrate statistical superiority of one dose over another. The objective is adequate characterization of each arm, not a hypothesis test. Sponsors who have read it as requiring a fully powered head-to-head dose comparison may have talked themselves out of doing it at all.

Guidance does not cover first-in-human starting doses, and it does not specifically address cell and gene therapy products. The European Medicines Agency (EMA) had not issued an equivalent guideline at the time of writing, and its expectations are currently expressed through scientific advice and the Cancer Medicines Forum, where early adoption of dose optimization principles has been described as essential.[7]

Feasibility and protocol conduct

This source of uncertainty offers wide latitude to be engineered out of a clinical study. It's also one of the hardest to specifically quantify reliably. Familiar claims that a fifth of sites enroll nobody or that four in five trials miss enrollment timelines are more than a quarter-century old and nearly impossible to trace to primary sources.

Better-documented numbers are less dramatic, more useful, and somewhat more telling of the state of the industry on this point. A study of Phase 2/3 interventional trials registered as closed in 2011 found 19% either terminated for failed enrollment or completed with under 85% of planned enrollment. This meant that roughly 48,000 participants did enroll in trials that could not meaningfully answer their primary question.[8] Tufts Center for the Study of Drug Development data, an industry dataset rather than peer-reviewed research, covering nearly 16,000 sites across 151 studies found 11% of activated sites enrolled zero patients, 37% under-enrolled, and 53% had to extend enrollment timelines.[9,10] The fact that these datasets are now over a decade old is itself a comment on how seriously the industry takes its most common operational failure.

In a study covering 16,000 sites across 151 studies

11%	enrolled zero patients
37%	under-enrolled
53%	extended enrollment timelines

What makes this source of uncertainty so distinctive is that the debt it assigns to clinical program capital is visible as it accrues. Enrollment shortfall is observable in real time. And yet, this uncertainty is still routinely absorbed as a schedule problem rather than recognized as an interpretability problem. That is, until rescue sites are activated late, under pressure, with compressed training and divergent protocol conduct, and the heterogeneity this introduces shows up in the primary analysis.

Endpoint and estimand

Three axes are rarely optimized together: sensitivity to the therapeutic mechanism, acceptability to the regulator, and meaningfulness to the patient. The estimand is a clear example of uncertainty debt, because ICH E9(R1) was adopted by ICH in 2019, has since been implemented by FDA and EMA, and sets

out how to address it.[11] An estimand has five attributes (treatment, population, variable, intercurrent events, and population-level summary) and ICH E9(R1) names five strategies for handling intercurrent events (e.g. discontinuation of assigned treatment, use of an additional or alternative treatment, and terminal events such as death): treatment policy, hypothetical, composite variable, while-on-treatment, and principal stratum.

A clinical program that specifies these strategies prospectively largely eliminates the uncertainty debt associated here. A program that does not will find that after unblinding, what was estimated becomes a matter of opinion. And it's important to note that this is a source of uncertainty where the debt can be repaid in public, in front of an advisory committee.

Uncertainty changes shape by modality

Uncertainty takes a different form in every modality. More importantly, the impact differs enormously from one modality to the next, because they differ in how irreversible they are. And as stated earlier, irreversibility is what makes uncertainty expensive. Three questions are constant across all modalities: what is the dominant form of ambiguity here, what does it cost, and when is it no longer able to be mitigated? Table 2 demonstrates increasing levels of irreversibility of uncertainty as it applies to five critical therapeutic modalities.

What holds for uncertainty across modalities is that as modality complexity rises, irreversibility rises, per patient, per batch, and per program. The same unresolved question of uncertainty costs more where you cannot redose, cannot re-manufacture identically, and cannot reverse the intervention.

As such, a posture of deferral, "we will characterize that in Phase 3," is not equally available across modalities. In small molecules it can be a reasonable trade. In cell and gene therapy it is rarely available, and treating it as available can be a costly error.

Table 2: Effects of uncertainty debt for five therapeutic modalities

Modality	Dominant form of uncertainty	Irreversibility	Where the effects can be mitigated
Small molecule	Exposure–response. Inactive, or under-exposed?	Low. Dose can be changed, formulation reworked, studies repeated at moderate cost	PK/PD and target engagement early; more than one dose into the pivotal
Antibody and protein	Immunogenicity and target-mediated disposition confounding exposure; tissue penetration	Moderate. Process changes demand comparability; anti-drug antibodies cannot be designed away retrospectively	Validated ADA assay from first-in-human; exposure analyzed by immunogenicity status; PD marker in target tissue
Autologous cell therapy	The product differs for every patient. Was it the product, the patient, or the disease burden?	High. One product per patient; manufacturing changes mid-program create comparability cliffs	Potency assay qualified against clinical outcome as early as possible; product attributes captured per batch as covariates; process locked before the pivotal
AAV gene therapy	Durability of expression, and pre-existing immunity. Is waning effect dose, immune response, or biology?	Very high. Redosing is generally not feasible in current clinical practice because of anti-capsid immunity	Expression biomarker with a pre-specified durability hypothesis; NAb screening strategy defined up front; long-term follow-up architecture designed at IND
In-vivo gene editing	On- and off-target activity in the tissue that matters; absence of evidence versus evidence of absence	Absolute. The edit is permanent, per patient	Off-target assessment strategy agreed with regulators before first-in-human; tissue-level editing readout; long-term safety architecture priced into the plan

Designing backwards from the label and the price

Sponsors often build strong regulatory strategies while giving payers far less attention. Good drugs can stall after approval for reasons that were fully determined years earlier at a protocol design meeting where nobody in the room was representing the payers' concerns. The asymmetry in concerns is easy to see. Regulators ask whether a product is safe and effective. Payers ask whether it is better than what they are already funding, by how much, and for whom. I have seen this pattern too many times: a sponsor secures regulatory approval without a pricing, market access, and reimbursement strategy in place, and the time gained in development is lost after approval due to the lack of payer strategy.

Regulatory and payer questions require very different evidence, and both can be answered by the same design decisions that are, unfortunately, often made with only one set of considerations in mind. An ambiguous value proposition (i.e. payer considerations) is the same failure as an ambiguous efficacy readout. A placebo-controlled trial, for example, can satisfy a regulator completely and still be unable to answer the payer's questions due to the lack of a comparator study. This form of uncertainty is costly for any clinical program, because it surfaces late in the develop/test/approval timeline.

Strategies I would suggest for the engineered alternative are not complicated. Extend the target product profile into a target value profile. Answer the question, "What must be true at launch for this to be reimbursed at the price the model assumes?" Seek payer advice in parallel with regulatory advice, rather than after it. Decide the comparator strategy once, for both audiences, before Phase 3 locks. And collect patient-relevant and health-economic evidence prospectively inside the study, because it cannot be added after the final readout.

While it may sound cynical to some, the reality is that the label is not the only finish line for clinical programs. The price is just as important.

Identifying and retiring uncertainty debt before the next capital gate

Optimism in clinical studies is structural

The numbers going into clinical study plans are not always entirely honest. Not because anyone is being dishonest, far from it. Recurring estimation failures are the result of competent people supplying optimistic inputs from positions where that optimism is rational. For example, feasibility questionnaires answered by investigators who want the trial at their healthcare network. This is the momentum of funding being sunk and lost at a gate where study continuation has a champion, but stopping has none. These estimation failures recur even in well-run companies, and the remedy is structural. Judgement should be relocated away from the party with an interest in the answer. Two devices can help discover unlooked-for liabilities, instead of simply managing it.

> As the trial is in the design phase, make this request of the entire leadership team: "It is three years from now. This program has failed badly. Tell me why."

The Red Team — This is not a reviewer, but an assigned adversary from outside the program, given full access and a specific task to find the most credible path by which this program fails, and the evidence required to rule it out. An assigned adversary has permission that a colleague does not.

The Pre-mortem — As the trial is in the design phase, make this request of the entire leadership team: "It is three years from now. This program has failed badly. Tell me why." It outperforms asking what the risks are, because explaining a fixed outcome engages different reasoning than forecasting an uncertain one.

A framework for scoring uncertainty

In Table 3, we present a structured diagnostic framework across nine clinical program domains (mirrored earlier in Table 1) and organized around a single question: if this reads out negative, will we be

able to identify why? Here, I want to emphasize what this framework is, and what it patently is not. The framework provided here is a structured diagnostic and conversation prompt, not a validated instrument, and it does not produce a validated composite score. It can also be more effective if rated by someone outside the program.

Each domain receives one of four statuses. Green means the question has been asked and answered decisively. Yellow is where uncertainty debt often hides. Red identifies a recognized liability, which is usually already known and already being discussed. Not assessed means the program has not formally asked or answered the question. This framework helps identify uncertainty that may not yet be visible to the program. The output isn't a grade, but a list of potential sources of uncertainty currently embedded in the program and an estimate of when the effects will be felt.

Table 3: Uncertainty scoring framework

Domain	Green	Yellow	Red
Target and translation	Human causality evidence; engagement readout qualified and fielded	Mechanistic rationale; engagement measurable but unvalidated	Animal efficacy only; no engagement readout planned
Product, CMC, and comparability	Potency assay qualified against clinical activity; product attributes characterized; process locked before the pivotal	Potency assay in development; partial characterization; comparability strategy outlined	No qualified potency assay; process changes planned mid-program without a comparability strategy
Dose, exposure, and pharmacology	Exposure-response characterized for efficacy and toxicity; more than one dose forward	Exposure-response modeled; single dose forward with stated justification	MTD carried forward; no exposure-response work
Endpoint, estimand, and statistical inference	Endpoint position documented with regulators; estimand and intercurrent-event strategy specified; assurance calculated from a discounted prior	Endpoint conventional; estimand partially specified; sample size based on power alone	Endpoint chosen for sensitivity alone; intercurrent events unaddressed; optimistic effect size with no sensitivity analysis
Population and enrichment	Prospective enrichment; criteria priced against eligible pool; placebo response mitigated	Criteria conventional; eligible pool estimated	Broad criteria; no enrichment rationale; no placebo-response plan
Feasibility and protocol conduct	Sites selected on performance history; real-world eligibility funnel; rescue triggers pre-defined	Mixed history and survey; funnel estimated	Feasibility questionnaires only; no competitive mapping
Safety	Mechanism-based hazards enumerated with pre-specified monitoring; background rates established; program-level aggregate detection	Hazards enumerated; monitoring partly pre-specified	Regulatory-minimum plan; no background rates; per-study detection only
Regulatory alignment	Positions on endpoint, population and analysis documented in minutes	Advice sought; some positions inferred	Alignment assumed
Value and access	Target value profile defined; comparator serves both audiences; health-economic data embedded	Payer input sought late; comparator adequate for regulator only	No payer strategy before Phase 3 lock

Note: Any domain may also be rated "Not assessed" when the program has not formally asked or answered the question. Red identifies a recognized liability, while "Not assessed" identifies uncertainty that may not yet be visible.

Conclusion

Uncertainty debt in clinical programs is cheap to accrue, invisible, and repaid late and at a high price. It is the mechanism by which reasonable, early, defensible decisions become uninterpretable results years later. Biology determines whether a program fails, but uncertainty determines what the failure costs.

De-risked clinical programs are not optimizing for success. They are optimizing for decisiveness and designing every study to produce an answer that can be acted on, front-loading the cheap questions before the expensive ones emerge. In my experience, success, when it comes, is largely a byproduct of not having spent years chasing ambiguous challenges.

Before the next capital gate, every development team and board should answer three questions. What uncertainty remains? Which of it is avoidable? And if the next study is negative, will we know why?

About the author

Panteli Theocharous, Ph.D., FRCPath, Chief Therapeutics and Clinical Strategy Officer, President of Europe

Dr. Panteli Theocharous serves as Chief Therapeutics & Clinical Strategy Officer and President of Europe at Worldwide Clinical Trials, where he leads enterprise strategy and drives growth in complex therapeutic areas. With over 35 years of executive leadership spanning CRO, pharmaceutical, and biotechnology environments, he brings extensive expertise across the full drug development continuum. His early clinical background in stem cell transplantation laid the foundation for recognized contributions in hematology, oncology, rare disease, and cell and gene therapy. Dr. Theocharous has held key senior roles across the industry, including Vice President and Global Biopharma Chief Medical Officer at PPD, as well as CMO and CEO positions at biotech organizations. He holds a Ph.D. in Clinical Cancer Medicine from University College London, is a Fellow of the Royal College of Pathologists, and earned degrees in hematology and biochemistry from King's College London.

References

1. Hay M, Thomas DW, Craighead JL, Economides C, Rosenthal J. Clinical development success rates for investigational drugs. *Nat Biotechnol*. 2014;32(1):40-51. doi:10.1038/nbt.2786
2. Citeline. Why are clinical development success rates falling? *In Vivo*. Published online May 16, 2024. Accessed October 2, 2026. https://invivo.citeline.com/IV154612/Why-Are-Clinical-Development-Success-Rates-Falling
3. Harrison RK. Phase II and phase III failures: 2013-2015. *Nat Rev Drug Discov*. 2016;15(12):817-818. doi:10.1038/nrd.2016.184
4. Bowling H, Cocucci A, Koo DCE, Harrison RK. Analysis of phase II and phase III clinical trial terminations from 2013 to 2023. *Nat Rev Drug Discov*. 2026;25(5):345-346. doi:10.1038/d41573-025-00208-6
5. Gupta KK, Attri JP, Singh A, Kaur H, Kaur G. Basic concepts for sample size calculation: Critical step for any clinical trials!. *Saudi J Anaesth*. 2016;10(3):328-331. doi:10.4103/1658-354X.174918
6. US Food and Drug Administration. *Optimizing the Dosage of Human Prescription Drugs and Biological Products for the Treatment of Oncologic Diseases: Guidance for Industry*. August 2024. Accessed October 2, 2026. https://www.fda.gov/media/164555/download
7. European Medicines Agency. Cancer Medicines Forum, December 2024. Accessed October 2, 2026. https://www.ema.europa.eu/en/events/cancer-medicines-forum-december-2024
8. Carlisle B, Kimmelman J, Ramsay T, MacKinnon N. Unsuccessful trial accrual and human subjects protections: an empirical analysis of recently closed trials. *Clin Trials*. 2015;12(1):77-83. doi:10.1177/1740774514558307
9. Tufts Center for the Study of Drug Development. 89% of trials meet enrollment, but timelines slip, half of sites under-enroll. *Tufts CSDD Impact Report*. 2013;15(1).
10. Getz KA. Enrollment Performance: Weighing the "Facts." *Applied Clinical Trials*. 2012;21(5). Accessed October 2, 2026. https://www.appliedclinicaltrialsonline.com/view/enrollment-performance-weighing-facts
11. US Food and Drug Administration. *E9(R1) Statistical Principles for Clinical Trials: Addendum: Estimands and Sensitivity Analysis in Clinical Trials. Guidance for Industry*. May 2021. Accessed October 2, 2026. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/e9r1-statistical-principles-clinical-trials-addendum-estimands-and-sensitivity-analysis-clinical

Worldwide Clinical Trials

About Worldwide Clinical Trials

Worldwide Clinical Trials (Worldwide) is a leading full-service global contract research organization (CRO) that works in partnership with biotechnology and pharmaceutical companies to create customized solutions that advance new medications – from discovery to reality.

Anchored in our company's scientific heritage, we are therapeutically focused on cardiovascular, metabolic, neuroscience, oncology, and rare diseases. Our deep therapeutic knowledge enables us to develop flexible plans and quickly solve problems for our customers.

For more information on Worldwide, visit www.worldwide.com or connect with us on LinkedIn.