A Horizon Europe 2028-34 FP10 Budget Question: What Happens When Discovery Gets Cheaper and Faster?

Prepared by Patrick Halford, Gaoithe Advisory – June 2026

Why Europe's next research and innovation programme needs to link budget logic to discovery compression, commercialisation speed and continuous feedback loops

FP10 is the next EU Framework Programme for Research and Innovation.

In current European Commission language, it is the proposed Horizon Europe programme for 2028-2034: a large research and innovation programme intended to strengthen Europe's science base, competitiveness, innovation capacity and ability to respond to real-world challenges.

That is the official frame. This note is about the pressure underneath that frame.

1. A missing question

Is Europe designing FP10 around the old research model - slow discovery, long applications, fixed consortia, fixed budgets and late exploitation - just as AI starts to make discovery faster, cheaper and more continuous?

That is the issue. Not the budget debate on its own.

The budget debate only makes sense if it is tied to the new speed of discovery and commercialisation. If discovery accelerates, the economics of FP10 change. If commercial routes appear earlier, the economics of FP10 change. If AI reduces the cost of research tasks during a project, the economics of each consortium change. If a project's assumptions can be monitored continuously, the economics of scope, budget and delivery change. If funded science starts producing spinout signals earlier, the economics of exploitation change.

So the question is not only how large FP10 should be.

The sharper question is what happens to the shape, timing and use of FP10 money when the cost and speed of discovery begin to move. That is where the current debate is too thin.

More ambition is not enough. More calls are not enough. More consortia are not enough. More reporting is not enough. FP10 has to be designed for a world where the cost of learning is no longer fixed.

2. Discovery is entering a compression phase

AI is not just another technology theme competing for programme funding. It is beginning to change the research process itself.

It is compressing the time needed to review evidence. It is reducing the cost of testing options. It is widening the number of hypotheses that can be explored. It is making simulation, modelling and experiment design faster. It is exposing commercial routes and spinout possibilities earlier. It is turning project intelligence from a retrospective reporting task into something closer to a live operating layer.

That changes the baseline for FP10. It changes what one euro should be able to do.

It changes how calls should be designed, how consortia should be built, how project scope should be managed, how budgets should be justified, how commercialisation should be tracked, when spinout potential becomes visible, and how fast assumptions should be tested again.

A proposal written in 2027 may be evaluated in 2028 and delivered into 2029 or 2030.

By then, the tools available to the consortium may have changed. The global competitors may have moved. The market may have shifted. The industrial need may have sharpened. The dependency risk may have become more urgent.

The original project logic may still be valid. Or it may be too slow. Too narrow. Too expensive. Too manual. Too late. That is why FP10 cannot be designed as a bigger version of the old machine.

3. The budget should follow the acceleration logic

This is where the FP10 budget discussion needs to move.

The serious question is not simply whether Europe spends more. **The serious question is what the budget is now buying.**

If AI compresses discovery, every budget request becomes less stable. Not weaker. Less stable.

A consortium may request funding based on one estimate of labour, testing, modelling, coordination, data analysis, market analysis and exploitation effort. Then AI changes the cost of some of those activities.

A work package that looked like twelve months of manual work may become six months of AI-supported work. A market-mapping task that once needed an external study may become a live evidence layer. A literature review that took weeks may become continuous. A technical pathway that looked uncertain may become clearer through simulation. A commercial route that was expected after the project may appear in year one. A spinout signal may appear before the midpoint.

So the initial budget is no longer just a funding request. **It is an opening hypothesis.** FP10 should treat it that way.

The budget should be tested. The scope should be tested. The partner logic should be tested. The exploitation route should be tested. The project economics should be tested again as the science moves.

This does not mean chaos. It means discipline. A serious FP10 project should not be locked into old budget assumptions when the work itself has changed.

If AI reduces the cost of one activity, the saved capacity should not vanish into process. It should move toward higher-value output: more testing, more validation, more deployment, more industrial engagement, more spinout preparation, more infrastructure, more capability and more useful science.

That is the budget issue. Not less money for Europe. **Better movement of money through the system**.

4. If you can fold a protein, you can fold a budget request

There is a practical question FP10 should put near the centre of its design.

If AI can help discover new materials, drugs, chemicals, molecules, devices and industrial processes, why would Europe not use similar tools to simulate and stress-test research budgets?

Put it bluntly: if you can fold a protein, you can definitely fold a budget request.

That does not mean replacing political judgement. It does not mean replacing expert reviewers. It does not mean handing the programme to a machine. It means giving human judgement better prediction machinery.

At the level of an individual consortium, AI could help test whether the project scope, partner mix, budget, work packages, exploitation route and technology assumptions still make sense as the science moves. It could ask hard questions before money is committed.

Is the labour estimate still credible? Is the modelling budget too high or too low? Is the partner mix right for the current state of the technology? Are three partners being funded to do work that one shared tool could now support? Are coordination costs justified? Is the exploitation budget too late and too small? Is the spinout pathway visible earlier than the project plan admits? Is the consortium still priced around manual work that AI can now compress?

At the level of the whole FP10 programme, AI could help simulate budget allocation, portfolio gaps, duplication risk, dependency exposure, industrial pull, spinout potential and the expected learning velocity of different funding routes.

That should become normal.

FP10 should not be governed only by annual calls, fixed templates and retrospective reporting. It should be governed by live intelligence. FP10 should use AI not only as a subject of funding.

It should use AI as part of the operating system for funding itself.

5. The old question and the better question

The old question was:

Can we win the call?

The better question is:

Should this call, consortium, budget and commercialisation logic still exist by the time the money moves?

That remains the core FP10 clock problem. But there is now a second layer.

If AI compresses discovery, project assumptions do not only drift because the outside world moves. They also drift because the tools inside the project improve.

A project may start with one view of cost. Then AI changes the work. A project may start with one view of scope. Then the science opens another route. A project may start with one exploitation plan. Then the market reveals a better one. A project may start with one consortium structure. Then the delivery need changes. A project may start with one budget logic. Then the economics of the work move underneath it.

The original plan may still be useful. Or it may be too slow, too narrow, too expensive, too manual or too late. FP10 needs to be designed for that reality from the start.

The question is not just whether Europe has enough money for research. The question is whether Europe knows how research money should behave when discovery accelerates.

6. Continuous feedback loops should shape budget decisions

FP10 should not treat budget as a one-time decision made at the front of the project. That is the old model.

A consortium applies. The budget is fixed. The work packages are fixed. The project begins. The world moves. The consortium reports. The reporting explains what changed after the fact.

That is too slow. FP10 should build continuous feedback loops into the economics of the project.

At call level, the programme should test whether the call logic still matches the science, the market and the strategic need. At application level, AI-supported tools should test whether the proposed budget matches the real cost of the work. At consortium level, the project should test whether the partner structure still fits the delivery need. At work-package level, the project should test whether AI has changed the time, cost or method required. At exploitation level, the consortium should test whether commercial routes, industrial users or spinout options are appearing earlier than expected. At portfolio level, FP10 should test whether too many projects are crowding the same ground while other areas remain underfunded.

This is where budgetary logic becomes dynamic.

A serious programme should be able to say: this work stream now costs less than expected; this saved capacity should move to validation; this commercial route has opened early; this budget should support a spinout test; this consortium needs a different partner; this project is duplicating work elsewhere; this scope is too narrow; this objective has been overtaken; this dependency risk has become urgent.

This is not bureaucracy. **It is the removal of blind spending.**

7. The application machine will also be compressed

There is another part of the FP10 system that should not escape scrutiny. The application machine itself.

If AI can compress discovery, it can also compress the work around discovery funding.

That includes identifying calls, analysing fit, forming consortia, mapping partners, drafting applications, checking duplication, testing work-package logic, reviewing budgets, comparing against previous funded projects, tracking global science movement and reallocating effort when the evidence changes.

Fully automated or semi-automated applications are no longer a distant idea.

Neither are self-assembling consortia, agentic application analysis, automated partner-fit scoring, live evidence checks, continuous project benchmarking and portfolio-level reallocation when scientific or industrial assumptions change.

That points to a different FP10 funding environment.

Less money absorbed by application friction. Less time lost in proposal machinery. Less duplication across competing consortia. Less manual effort spent proving what live systems can already detect.

More money reaching research, testing, deployment, exploitation and spinout formation.

How much of the research euro is actually reaching discovery?

The current system carries a heavy layer of preparation costs, bid-writing costs, coordination costs, application fees and success-based advisory fees. In some cases, those costs are substantial before any science, engineering, testing or commercial validation has happened.

AI will not remove the need for judgement, strategy or high-quality consortium design. But it should dramatically reduce the cost of routine application labour.

If that does not happen, Europe risks using AI to make the research frontier faster while leaving the funding gateway trapped in the old cost structure. That would be a strange outcome.

FP10 should be designed so that more of the money moves toward output, not process.

It should not use twenty-first-century AI to accelerate discovery while preserving a twentieth-century application machine.

8. Commercialisation cannot remain a closing chapter

The same logic applies to commercialisation. Under the old model, exploitation often arrives too late.

The project gets funded. The consortium works through the plan. Near the end, someone writes the exploitation section, maps stakeholders, prepares a commercial story and hopes the market has not moved on. That model is too slow.

In an AI-compressed research environment, commercialisation signals can appear earlier.

A customer problem may sharpen during the project. A regulatory shift may create a new opening. A competitor may validate a market. A procurement signal may reveal demand. A defence or industrial dependency may turn a research output into something urgent. A spinout path may become visible in month nine, not month thirty-six.

FP10 should not force all of that into end-of-project thinking. Commercialisation should become part of the feedback loop. Not as pressure to commercialise everything. That would be crude.

Some science should remain science. Some projects should produce knowledge. Some should produce standards. Some should produce infrastructure. Some should produce policy capability. Some should produce industrial capacity. Some should produce spinouts.

The point is not to turn every project into a startup. The point is to stop pretending that exploitation can be handled as a late appendix.

If new science creates a commercial route earlier than expected, the project economics should respond. If a spinout opportunity appears, the budget should support validation. If the market closes, the project should be able to redirect. If AI reveals a stronger adjacent use case, the consortium should not be punished for seeing it.

FP10 should treat commercialisation as a live signal. Not a final report section.

9. Rebaselining should begin before FP10 starts

The run-up to FP10 should not be treated as a waiting room. It should become a live rebaselining period.

Europe should be testing proposed themes, instruments, budgets and consortium assumptions against moving evidence before the programme begins.

What has AI already compressed? Which research activities now cost less than assumed? Which areas are moving faster than the expected call cycle? Which parts of project coordination can now be automated or semi-automated? Which industrial problems are becoming more urgent because of global dependency loops?

Which existing Horizon Europe projects are showing signs that their original scope is too narrow? Which projects are generating spinout potential earlier than expected? Which future FP10 areas need smaller, faster and more adaptive routes rather than large static projects?

Which project types are still priced around old labour assumptions? Which exploitation routes can now be tested before a proposal is even submitted? Which budget assumptions are already stale?

This is not bureaucracy. It is operating discipline. **A programme designed for a moving world needs moving assumptions.** If FP10 is designed against a frozen view of science, technology, budgets and commercialisation, it will enter the field already behind the play.

10. Technology clocks are still drifting

The previous FP10 clock argument remains central.

Technology clocks drift. Funding clocks are slower. Application clocks are slower still. Project clocks can be slower again. Commercialisation clocks often move too late.

That is the gap. AI makes the gap more visible.

A call can be written around one view of the world. The science can move. The market can move. The technology can move. The cost of work can move. The commercial route can move. The supply chain can move. The geopolitical pressure can move.

The project can then start under assumptions that were already weakened before the first meeting. That is not failure by the project team. It is a design flaw in the system.

FP10 should be built to see drift earlier. Not after review. Not after reporting. Not after the final exploitation workshop. Earlier. During call design. During application preparation. During evaluation. During grant agreement. During delivery. During scope review. During budget review. During exploitation and spinout assessment.

The funding system should learn while the funded work learns.

11. Scope and budget should not be frozen too early

One of the weakest habits in large research funding is the early freezing of scope and budget.

A consortium writes a project around a set of assumptions. Those assumptions become work packages. The work packages become budget lines. The budget lines become commitments. The commitments become reporting logic. The reporting logic becomes the project.

Then the world moves.

AI capabilities improve. A competitor releases something unexpected. A startup changes the commercial baseline. A geopolitical event changes the dependency map. A new model reduces the cost of a task by 80 percent. A customer problem becomes urgent. A scientific route opens that was not visible at proposal stage. A spinout route becomes more credible than the original exploitation plan.

Under old logic, this becomes a project-management problem. Under FP10 logic, it should become a design feature.

Funded projects should have structured ways to change scope and budget when the evidence justifies it. Not vague flexibility. Not chaos. Not every project chasing whatever is fashionable. Disciplined rebaselining.

If the original assumption is no longer strong, the project should be able to say so. If the industrial route has shifted, the project should be able to adjust. If AI has made part of the work faster or cheaper, the saved capacity should be redirected, not buried. If a spinout route appears early, the project should not have to wait until the end to take it seriously.

A research programme that cannot adjust to new evidence is not rigorous. It is rigid. And rigidity is expensive.

12. From funding competition to funding intelligence

FP10 should not only fund better projects. It should become better at knowing what should be funded, when, with whom, under what assumptions, and with what evolving budget logic.

That requires a shift from funding competition to funding intelligence.

The old model asks applicants to assemble the case, fight through the process and hope the assumptions still hold by the time the project starts.

The better model uses AI-supported intelligence to scan the whole surface: existing funded work, global science movement, industrial pressure, emerging competitors, spinout signals, policy dependency, partner capacity, duplication risk, budget productivity, technology clocks, commercial openings, delivery bottlenecks and cost compression.

That does not mean replacing evaluators. It means giving evaluators, programme teams and applicants a sharper picture of the world before money is committed.

It also means scope and budget should not be treated as sacred once the project begins.

If the science moves, the project should rebaseline. If the market opens earlier, the project should branch. If spinout potential becomes visible, the project should not wait for a closing exploitation package. If AI reduces the cost of one work stream, the saved capacity should move toward higher-value output.

If three consortia are unknowingly chasing the same ground, the system should see it. If one project creates reusable capability for another, the system should connect it. If an industrial dependency becomes urgent, the portfolio should respond.

That is not just better administration. That is better strategy.

13. The new measure: learning velocity per euro

FP10 needs better measures than activity volume.

The question should not only be how many projects Europe funds. It should be how quickly those projects learn, adapt, redirect and produce useful outputs.

A practical measure would be learning velocity per euro.

That does not mean reducing research to crude financial metrics. It means asking whether funded activity is increasing Europe's capacity to discover, decide, build and exploit faster than before.

Useful questions include:

- How much time did AI remove from literature review, modelling, simulation, experiment planning or reporting?
- How many project assumptions were tested during delivery, and how many were changed?
- How quickly did the consortium respond to technology movement?
- How early were exploitation routes tested?
- Did AI reveal adjacent markets, unexpected users or spinout opportunities?
- Did the project reduce European dependency in a measurable way?
- Did the work create reusable capability for other projects, firms, regions or sectors?
- How much of the budget reached output rather than process?
- How much application friction was removed?
- How much duplicated effort was avoided?
- How quickly did commercialisation signals appear?
- How much of the original budget logic changed once the work began?

That is a better conversation than budget size alone.

14. What FP10 should design for

FP10 should design for a different operating baseline.

1. Moving assumptions. Calls should be tested against live science, technology, market and dependency signals before they are published and during delivery.

2. Compressed applications. AI should reduce the cost, time and labour of proposal development, partner mapping, duplication checks and application analysis.

3. Adaptive consortia. Partner logic should be reviewed as technology and market conditions change. The right consortium at proposal stage may not be the right consortium eighteen months later.

4. Live scope and budget management. Projects should have disciplined ways to change direction and redirect budget when evidence justifies it.

5. Earlier commercialisation. Spinout, industrial and policy pathways should be tested while the project is alive, not after the useful window has narrowed.

6. Budget intelligence. AI-supported simulation should help programme teams understand where money is likely to create the highest learning velocity, reduce dependency, unlock capability or create reusable infrastructure.

7. Output over process. More of the FP10 euro should reach research, testing, deployment, exploitation and capability-building. Less should be trapped in avoidable application and coordination friction.

15. The uncomfortable truth

There is an uncomfortable truth here.

If AI compresses discovery but does not compress the funding process, then Europe will have modern tools trapped behind old gates. That is not a funding problem alone.

It is a design problem.

The research system cannot celebrate AI co-scientists in the lab and ignore AI intelligence in the funding model.

It cannot use AI to discover new materials but refuse to use AI to detect weak call assumptions. It cannot use AI to model molecules but avoid using AI to model portfolio risk. It cannot use AI to accelerate experiment design but keep project scope frozen around old assumptions. It cannot use AI to reduce discovery cost while tolerating unnecessary application cost. It cannot use AI to reveal commercial pathways while treating exploitation as an end-of-project ritual.

That gap will become visible. Fast.

The organisations that see it will move. The ones that do not will produce beautiful documents explaining why the world should have waited.

It will not wait.

16. The point

FP10 does not only need ambition.

It needs a new operating baseline.

AI is not just another technology theme to be funded. It is becoming part of the discovery machinery itself.

That changes the meaning of budget. It changes the meaning of project scope. It changes the meaning of consortium design. It changes the timing of spinouts. It changes the cost of applications. It changes the cost of learning. It changes the commercialisation clock. It changes how the programme itself should be governed.

The missing FP10 question is practical:

What should one euro of European research funding be able to do when discovery gets cheaper and faster?

If Europe cannot answer that, it risks funding the future through old machinery. And old machinery is not neutral. It has a cost.

It slows decisions. It protects friction. It hides duplication. It delays exploitation. It freezes budgets around assumptions that may no longer hold. It burns capacity before the real work starts.

FP10 has a chance to do something better. Not just a bigger programme.

A sharper one.

A programme that understands the new speed of discovery. A programme that uses AI to improve judgement. A programme that moves more money toward output. A programme that can rebaseline when the world changes. A programme that sees spinout potential early. A programme that treats funding as intelligence, not just allocation.

That is the real FP10 test. Not only how much Europe spends.

What Europe is now capable of learning, building and commercialising with every euro it spends.