A Philosophy of Research and Innovation Foresight: What the Estimation of European R&I Futures Assumes about Science

Veli Virmajoki
Finland Futures Research Centre, University of Turku

Abstract

Research and innovation (R&I) foresight estimates the possible futures of science, with method and with money. It has no philosophy. Philosophy of science, in the same two years, has produced most of the pieces such a philosophy would need, and pointed none of them at the field. Put together, the pieces give three results: the field runs on a positional ground of pursuit that no account of pursuitworthiness contains; it stacks three senses of readiness under one measure; and Popper's argument about the growth of knowledge limits what any of its estimates can claim.

Keywords: research and innovation foresight; philosophy of science; pursuitworthiness; readiness; horizon scanning; growth of knowledge

There is a practice that estimates the future of science, and it is not philosophy of science. It is done in Brussels, in Seville, in Luxembourg, in Paris, in The Hague and so on, by the analytical services of the Commission, by its agencies, by the consultancies it hires and by the think tanks that live around all this. In 2025 and 2026 it produced – I pick from what happens to be on my desk – scenarios of AI capability to 2030 (Hobbs et al., 2026), twenty-one key performance indicators for the digital economy to 2035 (Technopolis Group, 2026), scenarios of iridium demand to 2050 (Bolard et al., 2026), twenty-five "signals" of emerging technologies mined out of more than thirteen thousand proposals (EISMEA, 2026), a first monitor of research security (European Commission, 2026). What are these things? Not predictions, at least not in any simple sense. Nobody in these documents claims to see what science will be in 2035, that would be plain silly and the authors know it. They are estimates of what science and technology could do, should do and will be asked to do. Estimates made with method. Estimates acted upon with money – quite a lot of money, more than philosophy of science will see in its lifetime.

I have argued before that the estimating of possible futures of science is a topic philosophy of science has left aside (Virmajoki, 2022); that the criteria of pursuitworthiness philosophers discuss presuppose scenarios of the future of science that nobody has bothered to map (Virmajoki, 2023); and that philosophy, when it turns to a field concerning the future, should analyse the field through what it actually does – its methods, its concepts, the assumptions built in to them – and not through grand debates about the nature of time (Virmajoki, 2026). Here I point the three claims at one field, research and innovation foresight, *R&I foresight* for short. R&I foresight is the place where the futures of science are now estimated with method and with money. It has no philosophy. And philosophy of science has, during the last two years, produced most of the pieces such a philosophy would need – without pointing a single one of them at the field. Philosophers are busy people, I know. I am one.

Consider the two years. Philosophers of science have written on funding big science (Shaw, 2025), on lotteries for grant applications (Larue, 2026), on diversity of standards (Michelini and Osorio, 2026), on epistemic monocultures in AI-driven science (Fazelpour et al., 2026), on the

pursuitworthiness of experiments (Seifert, 2026), on what counts as novel when a machine learns something (Champion, 2025). Each of these analyses something R&I foresight does as a matter of routine. None of them mentions it. Then the other side. Thirty documents from the same two years bear on the futures of European research, and as far as I can see – and I have looked – not one of them refers to the philosophy of pursuit, of diversity, of evaluation, what have you. Two literatures, same years, same object, and they have never met. Why? The reason is not very interesting: different journals, different careers, different conferences, different coffee. What is interesting is what happens when they do meet.

Why Europe? Not because the argument is European. I discuss the European area as a case, for rather practical reasons. The practice writes itself down, estimates and decisions side by side, in documents anyone can read. And the comparisons with the United States and China, which are everywhere, bring the positional side of research into view. Whether foresight elsewhere holds the same assumptions is an empirical question, and I leave it open.

What makes a practice a field of inquiry? I take here, as elsewhere, the view of contemporary analytic philosophy of science (Virmajoki, 2026): a field is identified by how it actually operates – its methods, goals, objects, commitments – rather then by some definition from outside of what it ought to be. On that view R&I foresight qualifies, and rather easily. It has recurring methods: horizon scanning, expert-informed scenarios, composite indicators, technology readiness levels, benchmarking against named comparators, Delphi-type consultation and – this one is new – text-mining of the proposals that researchers write themselves (EISMEA, 2026; Santoleri et al., 2026). It has horizons, 2030, 2035, 2050, and notice that some of them are fixed by legal and budgetary cycles, not by anything in the phenomena (Technopolis Group, 2026, p. 30). Science is kind enough, it seems, to grow along the budget cycle. And it has a vocabulary that it uses consistently and almost never defines: *signal*, *trajectory*, *readiness*, *maturity*, *chokepoint*, *sovereignty*, *resilience*, and so on.

It has also a self-understanding. The EIC Tech Report says in so many words that inclusion in its list of signals implies no funding priority and exclusion no lack of relevance (EISMEA, 2026, pp. 8–9). The OECD paper assigns no probabilities to its scenarios and calls scaling laws consistent trends observed from past data, not immutable rules (Hobbs et al., 2026, pp. 10–13). The Digital Decade study says of its own targets that they are political, not calculated (Technopolis Group, 2026, p. 36). These are good methodological remarks – better, to be blunt, than some of what passes for methodology in futures studies literature. What the field does not have is a place where such remarks are collected, put side by side and asked what they add up to. That is what a philosophy of a field is for. One of the things, at least.

How, then, does R&I foresight operate in practice? How does foresight recognize what is emerging? Horizon scanning is the practice closest to what a futures studies researcher would recognise as foresight. The EIC Tech Report mined more than 13,380 proposals from 2021 to 2025, clustered them with the JRC’s text-mining tools, and got 411 candidate signals, which nine expert panels cut down to 25 (EISMEA, 2026, pp. 45–47). The criteria were novelty and relevance to EU policy (p. 47).

What does scanning assume? Ontologically, that technologies have trajectories – that an “early observable indication” indicates something, a path that is there, in some sense of there, before anybody has walked it. Epistemically, that novelty can be judged in advance, and judged from a

document. Champion distinguishes five kinds of novelty – surprise, blindness reduction, the elimination of deep ignorance, conceptual change, and theory-independent learning from data – and argues that none of them is necessary for a high-impact outcome (Champion, 2025, pp. 2–3, 9). The EIC's novelty is none of these in particular. It is whatever survives a text-mining cluster and a panel. I do not say this as a criticism; panels are what we have. I say it as an observation: the central concept of scanning has not been engineered, and philosophy of science put five candidate meanings for it on the table in the same two years. Nobody picked them up. Moreover, the instrument has a problem of its own. By end-2024 about 40% of the text in Horizon Europe proposal abstracts read as modified by a language model, and textual complexity predicts success for the human-written proposals but not for the assisted ones (Santoleri et al., 2026, pp. 8, 12). A method that mines proposals is mining text whose relation to the research behind it changes under its feet. The machine scans what the machine wrote. Whether the signals change with it, nobody knows yet. I do not know either.

How does foresight turn possible futures into targets? Targeting is fixing a number to a date. The Digital Decade study proposes twenty-one indicators at moderate and high ambition to 2035: edge nodes from 7,451 to 20,000 or 25,000, European cloud use at 30 or 40%, and so on (Technopolis Group, 2026, pp. 39–48). (The 7,451 is a measurement. The 20,000 is a wish with a round number, which is how wishes usually look.) The hydrogen deep dive runs a base and an innovation scenario to 2050, and under the base one world demand for iridium reaches 227% of today's supply (Bolard et al., 2026, Table 17).

Targeting is also where the field is most candid about itself. The targets are political, not calculated; the Digital Decade study says so. The hydrogen scenarios rest on an assumed capacity factor of 70%, and at 25% the EU's own iridium demand exceeds the global supply of 2022 by 2035 (Bolard et al., 2026, pp. 43, 68). And a study of health funding finds that strategic programmes reshape the portfolios of funded projects more than what in the end gets published (Fajardo-Ortiz et al., 2026, pp. 735–737). So the field knows three things, at least. A target is a decision dressed as an estimate. An estimate rests on parameters. And reaching a target at the level of projects does not reach it at the level of knowledge. What the field does not have is an account why the third is so. Popper has one: the growth of knowledge cannot be predicted by anyone who does not already have the knowledge that will grow (Popper, 1957). If one could, one would not need the programme; publish the target as a result and save the money. A target for 2035 that depends on what will be known in 2035 is, to that extent, not a forecast at all. It is something else, and we need to ask what. The field's own candour, political not calculated, is the right response. The field just does not say why it is right.

What happens when Europe measures itself against others? Benchmarking is measuring Europe against a named other, and it is everywhere. Europe produces 18% of AI publications and 2.77% of AI patents, while China holds 69.7% of the patents (Technopolis Group, 2026, pp. 20–21).

What does benchmarking assume? That the unit of value is the share. This is an assumption about what research is for, and it is the one philosophy of science has not seen. Consider the accounts of pursuitworthiness of the same two years. Shaw distinguishes normal, revolutionary and socially oriented big science, each with its features and its bugs (Shaw, 2025, pp. 5–11). Seifert finds five senses in which experimentation is worth of pursuit: knowledge, education, industry, innovation, history (Seifert, 2026, p. 8). Larue assumes all the way through that selection aims at the research most likely to be the best (Larue, 2026, p. 259). In none of these is there a ground of pursuit of the form: because a rival pursues it. The closest is the Manhattan Project, which Shaw lists among the

socially oriented cases (Shaw, 2025, p. 6) – the one case on the list where the good was that the other side did not get it first, and Shaw does not analyse it as such. The share is positional. What matters is not what Europe will know; what matters is how much of what gets known Europe will hold. R&I foresight runs on this ground. The philosophy of pursuit does not contain it, not even as a possibility to be dismissed. I find this remarkable. Philosophers, it seems, imagine science as a competition with nature. The field imagines it as a competition with China.

And what happens when foresight tries to steer research? Steering is aligning what gets done. The Digital Decade study recommends at minimum coordinated alignment with the budget cycle and ideally an integrated steering framework (Technopolis Group, 2026, pp. 78–101); the space task force wants to move "from coordination to orchestration" (Des Dorides et al., 2026, p. 2). *Orchestration* – the word I like, I admit. An orchestra has a conductor, a score, and nobody improvising.

What does steering assume? Epistemically, that alignment is good for inquiry, or at least not bad for it. Here philosophy of science has formal results, and they cut in a direction the field has not considered. Michelini and Osorio show that a community whose members hold different but legitimate standards explore more and, over a long enough horizon, outperforms even the best-organised homogeneous community – but only under total sharing in a dense network; otherwise an approach one agent rejects and another would have valued is lost, and lost for good (Michelini and Osorio, 2026, pp. 365–369). Fazelpour and colleagues show that a non-personalised AI returning the globally best configuration helps only when the problem is highly decomposable and use is moderate, and that it reduces diversity by about a quarter on average (Fazelpour et al., 2026, pp. 9–12). Now, steering is a move toward what these models call homogeneity, and research security, even under the field's own formula "de-risk, not decouple" (European Commission, 2026, pp. 15–16), is a move toward sparser sharing, at least at some joints. The models do not say steering is wrong. They say it is a trade, and they say what gets traded – which is already more than the field says. Sørensen and colleagues argue for temporal autonomy, the very thing the field treats as a cost to be cut (Sørensen et al., 2026, pp. 1–2). Several documents in my pile want Europe to be faster; for one of them speed is the key success factor of scale-ups (Schrijvers et al., 2026). Nobody has asked weather these are the same question. I suspect they are, deep down.

What follows from these practices? The analysis so far has been descriptive. However, three things follow that I want to state as results. They are, I dare to say, the first instalment of the philosophy the field lacks, and stating them is what makes this a paper and not a programme. A programme would end here, with a call for more research.

First. R&I foresight operates with a positional ground of pursuit, and the philosophy of pursuitworthiness does not contain it. The consequence is not that the philosophers are wrong. The consequence is that their apparatus does not transfer. Shaw's bugs and features turn upside down: crowding-out, the main bug of normal big science, is the very aim of a Chips Act; uncertainty concerns the rival's next move and not a framework's blind spots. Elsewhere I have argued that no criterion of pursuitworthiness is pursuitworthy *simpliciter*, only against a set of possible futures (Virmajoki, 2023). The positional ground is the clearest case of this I know. It makes sense only in a future where others pursue, and it is silent about every future where they do not – which is to say, about exactly those futures where it would have been waste of money.

Second. R&I foresight stacks three senses of readiness and lets a measure of one stand for the others. The AI4Bio report finds what it calls a maturity paradox: AlphaFold is mature within its domain yet sits at TRL 5–6 (Ruiz Serra et al., 2026, pp. 35–38). The paradox is an artefact of the stacking, or so I think. There is readiness to produce knowledge, readiness to deploy, readiness to hold a share, and a TRL measures only the technical part of the second. I flag this here and develop it in a separate paper, since engineering the concept properly, in the sense I have defended elsewhere (Virmajoki, 2026), is more than a section's worth of work.

Third. Popper's argument sets a limit on what any estimate in R&I foresight can claim. The growth of knowledge cannot be forecast. Readiness to produce knowledge is a claim about the growth of knowledge, so it cannot be forecast either. Readiness to deploy depends on it. Readiness to hold a share depends on it twice, since the rival's growth of knowledge enters as well. What we get is a ladder of unpredictability, and the field's methods sit on different rungs: horizon scanning near the bottom, where the claim is only that something is observable now; targets near the top, where the claim is about what will be known. "Political, not calculated" is what a target looks like when it is honest about its rung. What the field does not have is the ladder, and without it the honesty stays local, a remark in one study. Then again, a remark in one study is how features of a practice usually begin.

What, then, does this analysis establish? I have argued that R&I foresight is a field of inquiry, that it has no philosophy, and that philosophy of science has most of the pieces one would need. Each of the three results is also, I notice, a reason why the futures of science are a special topic and not one application of foresight among others. A positional ground of pursuit depends on other people's science. A stacked concept of readiness depends on what science will produce. And the limit on forecasting the growth of knowledge is a limit on forecasting the one thing R&I foresight most wants to forecast. Much remains to be understood. I have treated thirty documents as one field and a handful of papers as one toolkit, and both decisions can be questioned. I have said nothing about whether the field's assumptions are good ones to hold, which is perhaps the harder question. What I hope to have shown is that the futures of European research are estimated carefully, at scale and with method, by people who have a philosophy whether they asked for one or not – and that it is worth, for them and for us, to find out what that philosophy is.

References

Bolard, J., Nohl, M., Pavel, C., Istrate, R., Eynard, U., Blagoeva, D., Dolci, F., & Kuzov, T. (2026). *Deep dive on critical raw materials for hydrogen technologies in the EU* (JRC148105). Publications Office of the European Union.

Champion, H. (2025). Strong novelty regained: High-impact outcomes of machine learning for science. *Synthese*, 206, 134.

Des Dorides, C., Rodríguez, A. G., Martínez, J., & Schade, C. (2026). *Building a European competitive edge in space: An agenda for action* (CEPS Task Force Report). Centre for European Policy Studies.

EISMEA. (2026). *European Innovation Council tech report 2026*. Publications Office of the European Union.

European Commission, Directorate-General for Research and Innovation. (2026). *Research Security Monitor 2025: Raising awareness and building resilience* (Staff Working Document). Publications Office of the European Union.

Fajardo-Ortiz, D., Thijs, B., Glänzel, W., & Sipido, K. R. (2026). Evolution of public funding for collaborative health research towards higher-level patient-oriented research: A comparison of the European Union Framework Programmes to the program funding by the United States National Institutes of Health. *Quantitative Science Studies*, 7, 717–742.

Fazelpour, S., O'Brien, J., & Rubin, H. (2026). Navigating epistemic monocultures in AI-driven science: A simulation study. *Philosophy of Science*. Advance online publication.

Hobbs, H., Docherty, D., Aranda, L., Sugimoto, K., Perset, K., & Kierzenkowski, R. (2026). *Exploring possible AI trajectories through 2030* (OECD Artificial Intelligence Papers No. 55). OECD Publishing.

Larue, L. (2026). Funding research randomly. *Journal of Applied Philosophy*, 43(1), 258–275.

Michelini, M., & Osorio, J. (2026). Can scientific communities benefit from a diversity of standards? *Philosophy of Science*, 93(2), 353–377.

Popper, K. R. (1957). *The poverty of historicism*. Routledge & Kegan Paul.

Ruiz Serra, V., Fortuna, P., Curion, F., Petrillo, M., Leoni, G., Bili, D., Melandri, D., Vangheluwe, N., Comte, V., Bertolini Agnoletto, L., Coecke, S., Ceresa, M., Querci, M., & Wiesenthal, T. (2026). *Artificial intelligence for biology: Capabilities, readiness, and policy implications* (JRC147549). Publications Office of the European Union.

Santoleri, P., Rentocchini, F., & Lelli, F. (2026). *LLM-assisted proposal writing in competitive R&D funding: Evidence from Horizon Europe* (JRC Working Papers on Territorial Modelling and Analysis No. 02/2026). European Commission.

Schrijvers, M., Zuurdeeg, V., & Stolwijk, C. (2026). *Speed as key success factor for scale-ups* (TNO-2026-17904). TNO Vector.

Seifert, V. A. (2026). The pursuitworthiness of chemical experimentation. *European Journal for Philosophy of Science*, 16, 51.

Shaw, J. (2025). Funding big science: Managing diversity, social responsibility, and limited resources. *European Journal for Philosophy of Science*, 15, 74.

Sørensen, M. P., Lambert, M., Ravn, T., & Vostal, F. (2026). Time, quality, and integrity: Temporal autonomy as a missing link in research assessment reform. *Science and Engineering Ethics*, 32, 16.

Technopolis Group. (2026). *Study supporting the review of the Digital Decade Policy Programme: Final report*. Publications Office of the European Union.

Virmajoki, V. (2022). Understanding futures of science: Connecting causal layered analysis and philosophy of science. *Journal of Futures Studies*, 27(2).

Virmajoki, V. (2023). Pursuitworthiness in the scheme of futures. *European Journal for Philosophy of Science*, 13, 7.

Virmajoki, V. (2026). The shining: Illuminating philosophy and futures studies. *Futures*, 179, 103812.