Ready for What? Conceptual Engineering of Readiness in Research and Innovation Foresight

Veli Virmajoki
Finland Futures Research Centre, University of Turku

Abstract

Readiness is everywhere in European research and innovation policy, and defined nowhere. I argue that readiness is a two-place notion – something is ready *for* something – and that the practice drops the second place. Then three things get stacked under one name: readiness to know, to deploy and to hold a share. The technology readiness level measures the technical part of the middle one. The stack rests on the first, which Popper's argument places beyond forecasting, so the most precise claims of the practice, its targets, rest on its least forecastable ground. The proposal is small: say for what.

Keywords: readiness; conceptual engineering; technology readiness level; pursuitworthiness; R&I foresight; growth of knowledge

Why does readiness need a closer look?

Readiness is a word that European research policy cannot do without. Technologies have readiness levels. Models in artificial intelligence are mature but not ready, or ready but not mature, depending which table you read. Industries are ready, or they are not, to cover 40% of what Europe will deploy in 2030. Research systems are made resilient – which is readiness in a security vest, or so it seems to me – and the whole Union is asked, since the Niinistö report, to be prepared (Technopolis Group, 2026, pp. 10–11). Prepared for what? That question is not in the documents. The word is everywhere. It is nowhere defined. (I counted: in the thirty policy documents I work from, some form of *ready*, *readiness* or *preparedness* appears in the great majority, and a definition in none.) I do not hold this against anybody; most working concepts are like that, and they work, more or less. I only want to ask what the word is doing, and whether it could do its job better if somebody looked at it for a moment – preferably before the next call text is written.

Beside the thirty policy documents there are fourteen papers on my table from philosophy of science, all from 2025 and 2026, on funding, pursuitworthiness, novelty and the like. The two piles do not cite each other. Not once, as far as I can see. They work on the same objects – what science may come to know, which lines are worth to pursue – and they do it without talking to each other. Readiness is the place where the silence costs something, or so I argue. Europe is a case here, and not because the argument is European. It is because the European practice writes its readiness down: in funding calls, in benchmarks with dates on them, in foresight with catalogue number. Conceptual machinery does not know geography. And by conceptual engineering I mean nothing grander than what I set out in an earlier paper (Virmajoki, 2026):

taking the concept as the field uses it, and asking whether the use could be made clearer by the field's own lights, not mine.

The result, to say it in advance, is simple, almost embarrassingly so. Readiness is always readiness *for* something. Take a bridge. Is the bridge ready? The question has no answer until we say for what: ready for walking across, yes; ready for a lorry, perhaps; ready for a column of tanks, no. Same bridge, three answers. *Readiness* is a two-place notion, then, and the second place does the work. The practice drops it. It speaks of readiness as if it were a property that a technology or a country just has, like mass, or like colour; as if a bridge could be ready the way it can be grey. And when the second place is dropped, three quite different things get stacked on top of each other. A new vaccine is ready to produce knowledge when the mechanism works in mice and nobody yet knows whether it works in people. It is ready to deploy when it is approved, manufactured and in the fridge of a health centre. It is ready to hold a share when Europe makes it, and not only imports it. Three readinesses, three kinds of evidence, three ways of failing. The practice has a measure for the middle one – for its technical part, to be exact – and where a number is wanted it is let to stand for the other two. Sometimes the field produces a result that it has to call a paradox, and I use that paradox as my way in.

How does the practice use readiness?

The technology readiness level is the field's instrument for readiness, and it is worth to recall how modest it was meant to be. Nine steps, built for engineering programmes that needed to know whether a component can be trusted on a mission. It measures how far an artefact has travelled from the laboratory towards operation. In its own house it does the job well. The trouble starts when it is invited to other houses, which happens often, since it is the only guest that brings a number.

Consider the JRC report on artificial intelligence for biology. It places its models on two scales, a three-tier maturity scale and the TRL, defines both operationally in two tables (Ruiz Serra et al., 2026, Tables 5–6), and finds that the two do not agree. AlphaFold is domain-mature – it is the standard, everybody uses it – and sits at TRL 5–6. OpenCRISPR-1 is nascent, at TRL 2–3 (pp. 35–38). The authors call this a maturity paradox and conclude that maturity is not the same as readiness (pp. 111–112). Right conclusion, honestly written down, which is more than most reports manage. But notice: it is a conclusion about a word. What the two things are, as kinds of readiness, the report does not say.

The other uses look like the same concept. They are not. The EIC Tech Report's signals show early scaling potential (EISMEA, 2026, p. 8): ready to grow. The clean-energy observatory reports coverage of 2030 deployment needs – battery cells at 21%, PV wafers at 0% (Georgakaki et al., 2026, Table 4): ready to supply what a target demands, a target somebody chose. The hydrogen competitiveness report finds most projects not bankable without public support (Houghton et al., 2026, p. 4): ready to run, not ready for a bank. And grid connections for AI data centres take up to ten years in Germany or the Netherlands (Bômont and Hobhouse, 2026, p. 2), so a model can be ready in every sense the AI4Bio report knows and still wait a decade for a

cable. Each use has a second place, and in each it is different. Where a number is wanted for an artefact, the TRL stands in – the EIC report attaches an estimated TRL to each of its signals (EISMEA, 2026, p. 46). Elsewhere every report builds a measure of its own – a coverage ratio, an index, a maturity model, and so on – and the measures does not talk to each other, because the word does not tell them that they measure different things.

Interstingly, the documents split readiness themselves, now and then. TNO separates technology readiness from manufacturing readiness (Schrijvers et al., 2026, p. 25). Each split is made locally and stays there; the next report starts again from the word. It is not a gap in what the field knows. It is a gap in what the field can say. (Institutional memory is not ready either. For what, I leave open.)

What can the philosophers tell us about readiness?

Philosophers of science hardly use the word *readiness*. They use the concept all the time, under nicer names – pursuitworthiness, promise, novelty – since philosophers like nicer names. Seifert argues that chemical experimentation is worth of pursuit in several senses, from knowledge production to industrial development, and that the same experiment can be worth pursuing in different senses at different times: spectroscopy was ready to discover elements in 1859, ready to determine molecular structures after quantum mechanics, and it is ready to train machine-learning models now (Seifert, 2026, pp. 8, 17–20). Larue defends a lottery for research funding, with a screen that removes only the applications failing minimal standards, because peer review is reliable at the top and at the bottom of the quality range and not in the middle (Larue, 2026, pp. 260–261). So there is a weak readiness, ready to be funded, which reviewers can judge, and a strong readiness to produce the best research, which they cannot reliably judge, whatever they think. (Panels think quite a lot.) And Fazelpour and colleagues say, in so many words, that AI preparedness is value-laden (Fazelpour et al., 2026, pp. 18–20). It is built, according to what one wants the technology for, and not read off from anything.

What do they share? Not a definition. Each keeps the second place and fills it: ready for knowledge, for industry; ready to be funded, weakly; prepared for whatever one values. None of these is the TRL, and none of them comes with a number, which is the one thing the practice most wants from a readiness concept. That is not a deficiency of the philosophers. It is a fact about the concept.

What, then, are the three objects stacked under readiness?

Let me now separate the objects. Under all these uses there are three, I think, and they come in an order.

The first is readiness to produce knowledge – *ready to know*. A programme or a model is ready in this sense when pursuing it would yield something not yet known. This is the philosophers' readiness, and it has the structure I argued for in 2023: it is always relative to a set of possible futures, and there is no readiness-to-know *simpliciter* (Virmajoki, 2023). Its characteristic failure

is the dead end. AlphaFold's domain maturity is readiness of this kind, already cashed in; the knowledge got produced.

The second is *ready to deploy*: an artefact can be used at scale, outside the laboratory, by people who did not build it. The TRL measures its technical part – not the manufacturing, not the regulator, not the bank. Its failure is the thing that works in the laboratory and not in the field. AlphaFold's TRL 5–6 is of this kind; a tool everybody uses in research is not thereby a tool for a clinic. There is no paradox in that. There are two readinesses.

The third is *ready to hold* a share. A country, an industry, a Union is ready when it will have meaningful part of a market, a technology, a supply chain – relative to others. Its failure is being late: the technology works, it is deployed, and somebody else owns it. Resilience is something else again, readiness to withstand and not a share of anything, and I leave it at the edge.

Now the order. You cannot deploy what nobody has found out, and you cannot hold a share of what is not deployed (a share of patents or of a standard can be held earlier, resting on the first directly). The stack rests on the first. The hydrogen deep dive projects world demand for iridium in electrolysers at 227% of today's global supply by 2050 under its base scenario and 52% under its innovation scenario (Bolard et al., 2026, Table 17, p. 53). The difference is not in how much hydrogen Europe wants. It is mostly in what will have been found out – whether catalyst loadings keep falling – and what gets built on it. The top of the stack swings by factor of four on what happens at the bottom.

The practice has one number for the second, the TRL. It has no number for the first, because nobody has one (the philosophers have criteria, which is different thing, and weaker). And it has many numbers for the third, because shares are easy to count. So it does what any practice with one good measure does: it uses the measure for what it measures, and then, when a number is wanted for the other two, uses it again. That is how the standard of a whole field comes to be "not ready". The maturity paradox is not a paradox. It is a stacking error – an instructive one, luckily for us, because the report caught it and wrote it down.

The error has a cousin running the other way. European heat-pump manufacturing covers 385% of 2024 deployment needs, splendid, but the figure reflects assembly and about 90% of compressors are imported (Georgakaki et al., 2026, p. 42). Ready to hold, on the level of the box. Somebody else is ready inside the box. To be sure, the practice is not confused; the AI4Bio authors know perfectly well that a TRL is not a judgement of scientific promise. It means that *readiness*, unindexed, has no place to put that knowledge. Conceptual engineering, in the modest sense I mean, is giving a practice such place. That is all. It is not nothing.

Why does the bottom of the stack matter?

Why does this matter, beyond tidiness? Because of what sits at the bottom.

Readiness to know is a claim about the growth of knowledge, and Popper's argument – which I take to be decisive, and I have argued elsewhere why – says that the growth of knowledge cannot

be predicted by anyone who does not already possess the knowledge that is going to grow (Popper, 1957). If we could say now what AlphaFold's successor will find out, we would not need the successor. Would save lot of money, that. This does not mean readiness to know cannot be judged. It is judged every day, by referees and panels, and sometimes they are right, which is worth to remember. But it is a judgement of promise against possible futures, and promise is not probability, however much we would like it to be. Then again, there is a new twist. By the end of 2024 about 40% of the text of Horizon Europe proposal abstracts read as modified by a language model, and textual complexity predicts success for human-written proposals and not for machine-assisted ones (Santoleri et al., 2026, pp. 8, 12). Readiness to know is judged from a text, and a text is a thing one can get better at writing without getting better at knowing. Now the machine writes it too. Whether the machine is ready, and for what, I leave to the reader.

Now trace the limit up. Readiness to deploy inherits whatever cannot be forecast in knowledge. Readiness to hold inherits it twice, since it depends on both, and on the rival's as well. The further up a claim sits, the more unforecastable material it is built out of – and, as it happens, the more decimals it comes with. Take European cloud use at 30% or 40% by 2035, depending which ambition is chosen (Technopolis Group, 2026, pp. 47–48): a share of a deployed technology, built on knowledge nobody has yet. The targets are the most precisely stated claims in the practice and they rest on its least forecastable ground. The Digital Decade study itself calls them political targets, not calculated ones (p. 36). Exactly right. That is what a claim at the top of the stack has to be.

The practice knows how to say this honestly, by the way. The OECD paper on AI trajectories builds four scenarios to 2030, assigns no probabilities, and advises to plan across the whole range (Hobbs et al., 2026, pp. 10–11, 41); the hydrogen deep dive gives its two iridium numbers side by side and does not pick one. The practice forgets this a little when the scenarios become one target, and the target a benchmark, and the benchmark, in the next report, a fact.

Ready for what, and by whose lights?

Who fills the second place, when it is kept? The practice answers this all the time without asking it. Consider Chinese open-weight models. They are free, good and run locally, and, as Bruegel points out, they create a new dependency for European firms (García-Herrero and Storella, 2026, p. 4). For the firm, ready to use. For the Union, ready to depend. Which readiness counts? The word does not tell. It cannot; it has no second place.

What would an engineered concept let us say?

The proposal is small. *Readiness* should be written with its second place – *ready to know*, *ready to deploy*, *ready to hold*. The TRL stays, as the measure of the technical part of the second; I have nothing against it. Readiness to know gets no number, because there is none to be had, but it gets what the philosophers have: criteria, and scenarios of the futures against which the criteria make sense. Readiness to hold keeps its benchmarks, with the dependency written in, so that a target for 2050 carries in its very statement the two unforecastable layers underneath.

What could the field then say? That an industry covering 21% of 2030 battery-cell needs is unready to supply its own deployment of a technology that is fully ready to deploy – a different problem, with different remedies, than a technology that does not work. And of its own targets, what the Digital Decade study already said once – political, not calculated – but as a feature of the method rather then as an admission in one report.

Much remains unclear, as usual. I have not shown that there are only three objects; readiness to grow may be a fourth, or the first under a different description, and I am honestly not sure which. And who should fill the second place I have not said; Fazelpour and colleagues are right to call it value-laden, and it is another paper, probably not mine. What I hope to have shown is narrower. The practice has a concept doing three jobs under one name. The philosophers have the tools to tell the jobs apart, and they have been using them, for two years, in papers the practice does not read. Readiness is always readiness for something. The least we can do is to say for what.

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.

Bômont, C., & Hobhouse, C. (2026). *Feeding the beast: Energy security as a precondition for the EU's AI power* (Brief 15). European Union Institute for Security Studies.

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

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

García-Herrero, A., & Storella, T. (2026). *China's artificial intelligence goals and strategic choices for Europe* (Analysis 24/2026). Bruegel.

Georgakaki, A., Piela, K., Delgado Callicó, L., Finck, R., Grabowska, M., Ince, E., Miletic, M., Mountraki, A., Nohl, L., Taylor, N., & Terca, G. (2026). *Overall strategic analysis of clean energy technology in the European Union* (CETO Status Report 2025; JRC146873). Publications Office of the European Union.

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.

Houghton, T., France, W., Long, J., Nkata, L., & Galeazzi, G. (2026). *Final report on European hydrogen competitiveness*. ERM for the Clean Hydrogen Partnership.

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

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, Joint Research Centre.

Schrijvers, M., Zuurdeeg, V., & Stolwijk, C. (2026). *Speed as key success factor for scale-ups: The Dutch scale-up ecosystem in international perspective* (TNO Report TNO-2026-17904). TNO Vector.

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

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

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.