CO2RE THE GREENHOUSE GAS REMOVAL HUB

Technology Learning for Net Zero

Deep uncertainty, two-factor learning and policy sequencing in carbon dioxide removal

October 2026

Siyu Feng, Diarmid Roberts, Joseph Stemmler and Mark Workman [1]

Summary

All Paris-aligned pathways for limiting warming to 1.5°C require large-scale deployment of carbon dioxide removals (CDR) by mid-century. Integrating CDR technologies into net-zero strategies is justified in the long term partly on the basis of economic efficiency: they can offset hard-to-abate and dispersed emissions where removal may ultimately be less costly than further mitigation. They may also enable a return of emissions to 1.5°C following an overshoot.[2] Technical removals are particularly important because of their potential scale and ability to lock up CO_2 on a geological timescale. However, the technical CDR sector is nascent at present, with the costs of bioenergy with carbon capture and storage (BECCS) at $400+/tCO_2$ and direct air capture (DAC) ranging from £400 to 2,500 /tCO_2. In the short term, this makes these technologies prohibitively expensive to integrate and scale consistent with Paris-aligned goals. Policies are therefore needed to support their development and rapidly reduce their costs towards the anticipated economic efficiency.

At present, policy analysis commonly estimates the learning rates of BECCS and DAC technologies based on single factor learning curves (SFLCs). These estimates are often calculated by benchmarking against analogous technologies and are subsequently incorporated into whole-energy systems models. These analyses of learning rates are fundamental for policy design as they shape estimates of future technology costs and, consequently, the scale, timing and form of targeted public policy interventions.

This policy briefing makes the case that orthodox approaches to learning curve analysis can blind policymakers to the complexity inherent in technology learning for three reasons: (1) they radically simplify the complexity of technology learning; (2) they implicitly assume that markets allocate resources when, in practice, firms do; and (3) they tend to apply SFLC assumptions exogenously, resulting in actors advocating opaquely for assumptions that they prefer. There is therefore a risk that poorly informed policies will be developed without the application of more sophisticated approaches. This will inflate the costs of achieving net zero ***and*** undermine one of the main justifications for bringing CDR technologies into the climate toolbox. We propose a 'deep uncertainty approach' to address the fundamental flaws in orthodox learning curve analysis. We seek to establish this process in the literature, and invite others to utilise it, adapt it and critique it.

1. Introduction

As part of its net-zero aspirations, the UK intends to establish a new carbon dioxide removal (CDR) sector from a standing start. Simulations suggest that the UK requires up to 100 $MtCO_2$ of CDR by 2050.[3] It has allocated more than £20bn indirectly and £100m directly to establish the sector since 2019.[4] It has also proposed policies designed to support the sector from the mid-2020s to 2030s – namely the integration of negative emissions credits into the UK Emissions Trading Scheme (ETS) and the introduction of a Carbon Contract for Difference (CCfD).[5] CO_2RE has undertaken analysis that suggests these policies are inadequate to facilitate the attraction of capital to the start-ups in the UK's nascent CDR sector, and that other risk sharing interventions are therefore required.[6] Furthermore, additional analysis has found that cost curve simulations for policy design have not been sufficiently taken into account, potentially leading to misinformed policy design for the UK CDR sector.[7] This has happened because of a failure of the analytical community across academia, consultancies and multilateral organisations to understand the characteristics of CDR, which we have argued is composed of both market-creation and market-led complexity.[8] We have also illustrated how these complexity characteristics are not well captured by existing tools that support policymaking to target CDR technology learning[9,10]

In this policy briefing, we provide an alternative approach to decision-making support which is both exploratory and participatory in nature to ensure technology learning is as effectively captured in analysis as is pragmatically possible. This framework allows multiple factors to be explicitly involved in learning, rather than embedded within opaque or speculative assumptions external to the model. ***We contend that learning dynamics are so poorly and inconsistently applied in the technology innovation corpus that there is a real risk of public policies failing to address the needs of technology developers or provide the public purse value for money.*** Learning curves are highly important because they shape the dynamic relationship between policy choices and technology cost trajectories over time. If analysis distorts insights, the policy built on it stands to be mis-calibrated – either requiring the provision of far more support than needed in reality, or too little to stimulate a meaningful cost reduction trajectory.

Using the UK as a case study, CO_2RE has designed and applied an analytical framework to gain insights on (1) the impact of policy interventions on firm-level strategic investment to generate multi-factor technological learning effects; (2) the quantified impact that this will have on CDR sector development; and (3) the fiscal implications of those interventions. The study employs an economic theoretic modelling approach that endogenises different learning curve impacts on technology development applied within an exploratory modelling[11] – a framework which in turn engages with market actors and policymakers to identify targeted interventions. This enables the analysis to reflect stakeholder market perspectives and generate insights relevant to policymakers designing UK CDR policy. The proposed approach challenges the dominance of consolidative,[12] whole-system optimisation approaches that focus on predicting an unknowable future and distort innovation policy design for net zero.[13] Because government research, development, demonstration and deployment budgets are constrained and the stakes of resource misallocation are high, more reliable targeting of innovation spending can directly improve investment decisions.[14] We seek to establish this process in the literature, and invite others to utilise it, adapt it and critique it.

Section 2 sets out learning-curve analysis, Section 3 describes how learning curves are presently compiled, Section 4 sets out what still needs to be considered, and Section 5 concludes and draws the analysis together. This briefing is targeted at a policy audience. For those that require a more substantive treatment of the theoretical economics that underlie the endogenisation of different learning curve impact on technology deployment see Feng, Stemmler and Workman (forthcoming).[15] For those that would like to better contextualise how this briefing augments the state of the learning curve and policy sequencing corpus in the assessment of technology deployment see Grubler and Wilson (2016).[16]

2. Learning curves in orthodox net-zero analysis

There is a rich literature on the role of learning curves[17] relevant to decarbonisation technologies,[18,19,20,21,22] including that specific to CDR.[23,24] The application of one-factor learning curves or single-factor learning curves (SFLCs) remains the most dominant method for quantitative modelling of energy technology innovation. This is due to the dominance of cost-optimisation models in policy design, which tend to only allow a single variable to be applied[252]. Gambhir, et al. identify three principal methods for projecting future technology costs, each of which have strengths and weaknesses: SFLC extrapolation, expert elicitation and bottom-up engineering assessment, alongside an increasingly common fourth practice of combining them.[26] In the case of a sector as nascent as CDR, it also means that learning rates are justified based on a very small corpus, often of other modelled analyses, which are laden with assumptions. An example of how a direct air capture (DAC) UK learning curve study was undertaken applying orthodox approaches is shown in Box 1.

SFLCs, however, capture only a subset of the drivers recognised in the innovation literature.[27] The omitted drivers are not minor. Technology innovation is a high-dimensional, dynamic, emergent and complex learning process encompassing the following components:[28]

1. Learning-by-researching, or -by-searching, in which cost reductions arise from technical improvements generated through research and knowledge creation.
2. Learning-by-deployment, which has the following sub-components:
   - 2.1 Learning-by-doing;
   - 1.2 Learning-by-using; and
   - 1.3 Learning-by-interacting (competition and x-efficiencies, which is similar to market demand), or spillover of learning, which has further granularity.[29]
3. Economies of scale at device, plant and industry level, which can also contribute to technology cost reductions, with device-level upscaling potentially having particularly significant positive or negative impact on costs.
4. Market factors:
   - 1.1 Production market drivers explain the dynamics related to the supply-chain, incorporating input material availability and their market price, the cost of supplies, the variation of capital and labour costs.
   - 1.2 Market demand accounts for the dynamics related to the sales market, encompassing an increased number of competitors and the implementation of new environmental regulations.

Quantitative, multi-factor learning curves methods (MFLCs),[30] and more recently, advanced versions of bottom-up cost models (BUCM),[31] are still in their infancy. This means that the other learning drivers, such as market dynamics and learning-by-interacting across different stakeholders and geographical areas, remain poorly quantified, despite their recognised impact on cost reduction.

Box 1: Orthodox approaches to learning curve development and technological diffusion for net zero technologies – in this case the cost of direct air capture (DAC) for the UK's 7th Carbon Budget (2038–2042)[32]

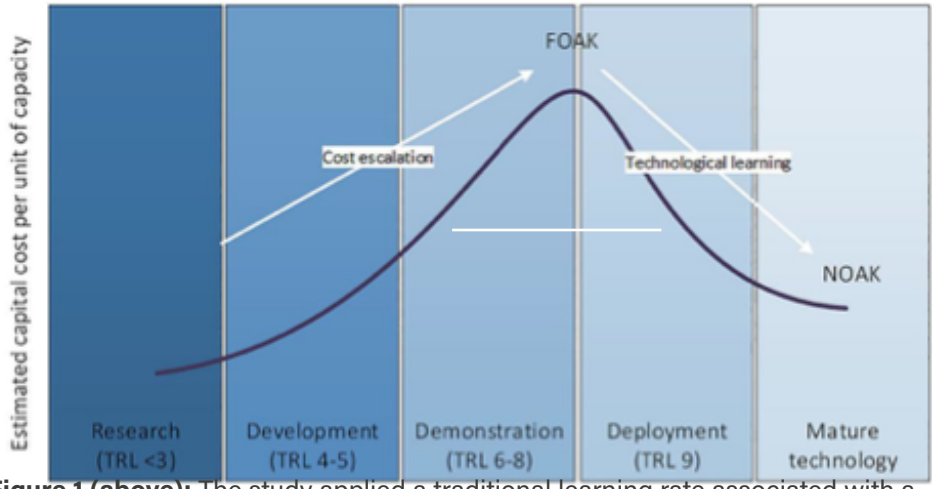
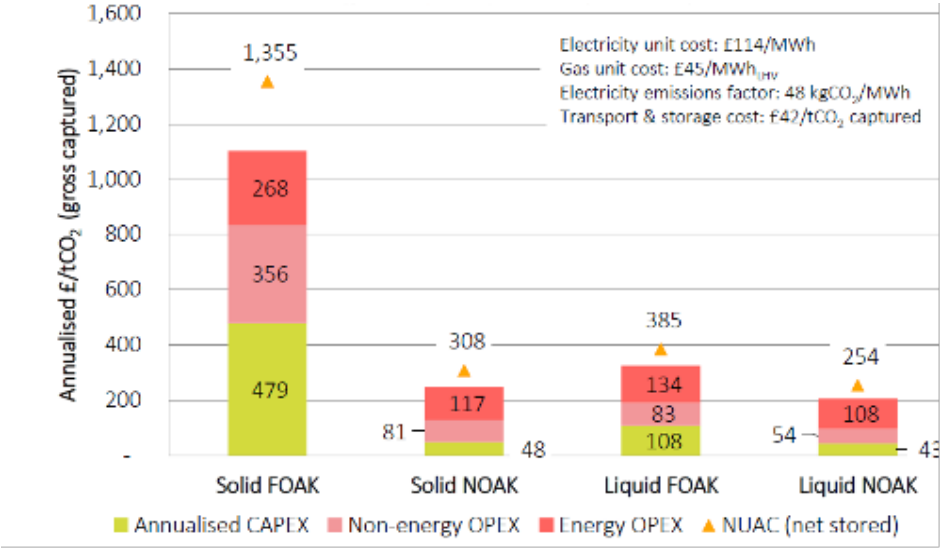


Figure 1 (above): The study applied a traditional learning rate associated with a first of a kind (FOAK), whereby costs increase and then, as learning takes place through scaling, costs fall to nth of a kind (NOAK).



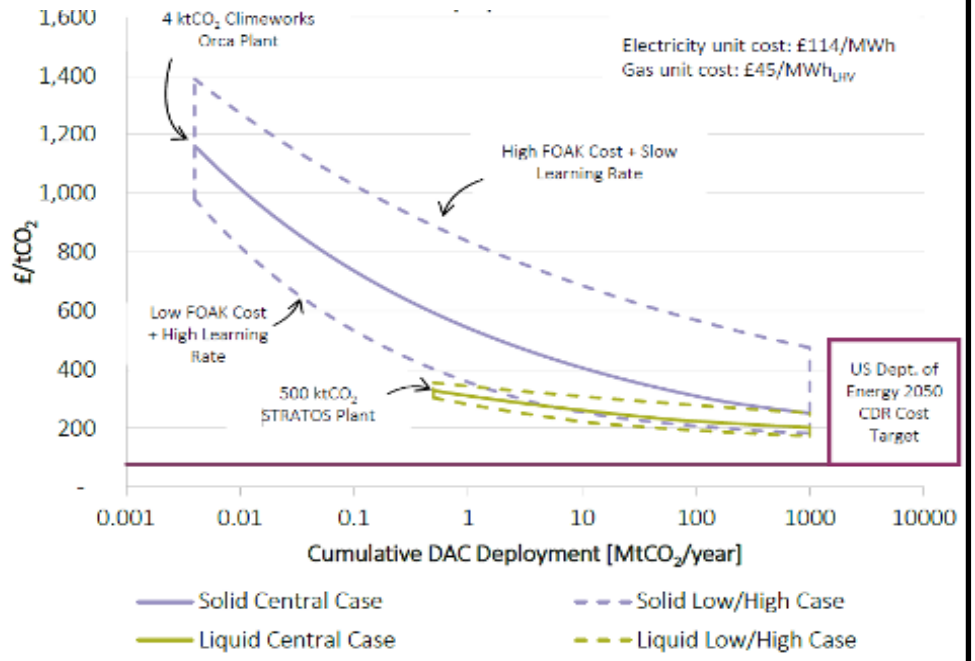


Figure 2 (above): A cost model was then developed based on academic literature for solid sorbent DAC (purple) and then liquid sorbent (green). Learning was based on historical growth and learning rates for analogous technologies, based on an arbitrary, globally modelled and scaled sector for both technologies, with 2050 costs simulated at £328/tCO_2 and £242 /tCO_2, respectively (central scenario). Uncertainty is addressed by having an upper bound learning rate and lower bound learning rate. These are then plotted on a graph **(Figure 3, left)**, with each of the costs broken down into CAPEX, non-energy OPEX and energy OPEX for each technology based on FOAK and NOAK.

This does not mean that SFLCs are without empirical value. Because cumulative deployment is often correlated with R&D, supply-chain development, market formation, standardisation and firm-level learning, SFLCs may indirectly capture some of these wider drivers. Two-factor learning curves (2FLCs) improve on this by adding knowledge-stock measures, such as cumulative R&D expenditure or patents. However, both remain reduced-form representations and cannot fully identify the policy, market, institutional and firm-level mechanisms through which different forms of learning are generated. This matters because each type of learning is stimulated by one or more policies, market interventions and firm responses to stimuli.

The emergent interaction of these interventions, and their direct and indirect impacts on a specific learning channel, therefore become difficult to disentangle. As a result, the effectiveness of an ecology of multiple innovation policies to stimulate individual and/or multiple learning rates is itself subject to deep uncertainty[4].

The oversimplification by which orthodox approaches assume away the complexity of technology learning has two further concomitant impacts as a function of their integration into whole-system optimisation modelling. These economic models implicitly suggest resources within an economy are allocated by market forces. In practice, they are allocated by governments through innovation initiatives such as the Inflation Reduction Act in the US, but most importantly via firms through strategic investment decisions. These will be predicated on a variety of objectives beyond least cost and profit maximisation, and are highly problematical to model.

Furthermore, the way in which whole-system optimisation models apply technology learning exogenously can encourage stakeholders to advocate for modelling assumptions that they know will lead to the strategy or policy that they already prefer. This could result in the distorting of policy design by actors advocating for policies that favour their specific technology at the expense of the economic efficiency of the sector as a whole.

3. A deep uncertainty approach to learning curves: a direct air capture use case

This study seeks to advance MFLC and BUCM approaches by developing a deep uncertainty framework that accommodates multiple factors in technology learning curves. By making them and their associated uncertainties transparent, the framework can support more targeted policy interventions aligned with different technology learning needs. It combines exploratory computer visualisation with a participatory and deliberative process involving relevant stakeholders. This enables stakeholders to explore how different assumptions, system components and policy choices interact, and to examine trade-offs across a wide range of plausible futures. The objective is therefore to identify interventions that are robust across uncertainty rather than optimal under a single assumed future.

As a specific case study, we opt to focus on DAC technology due to the simplicity of the 'value chain', whereby the study focus was on the removal technology component only. We engaged with several groups of individuals in the space to calibrate the application of our model to DAC. These expert-elicited inputs were then employed as parameters in a theoretical DAC model designed to evaluate the interactions between learning and policy-setting, including timing and duration. Among other surveyed items, these parameters include the individuals' perceptions of removal costs, learning rates and required levels of policy support. Both the model and the expert survey are described further below. Rather than trying to cover every possible form of learning as per above, the priority was to keep the analysis structured, focused and comparable.

The model applied:

- CDR cost reduction, through two main learning channels:
  - (1) ***Learning-by-doing*** – costs fall as firms deploy more capacity and gain production experience; and
  - (2) ***Learning-by-searching*** – costs fall through R&D investment and knowledge accumulation.
- We model learning-by-doing and learning-by-searching simultaneously, so that both the firm's production and R&D choices directly affect future costs.
- The dynamic model simulates a single CDR firm's removal output and investment/R&D responses to the allocation of resources relevant to innovation and to a sequence of policies over a 25-year period.
- The model maps policy instruments into removal prices faced by the firm over time (i.e. a CCfD strike price) – for detail, see Feng, Stemmler and Workman (forthcoming, endnote 15). Given these prices, the firm decides how much CDR to produce, and how much to invest in R&D. These decisions affect future costs through learning-by-doing and learning-by-searching.

In this way, different policy sequences can lead to different deployment pathways, investment patterns and long-term costs.

Because empirical data on emerging CDR technologies remains limited, we conducted a targeted expert elicitation exercise from UK CDR specialists to inform numerical examples.

This comprised a structured pre-workshop questionnaire followed by group discussion. Participants were asked to provide plausible ranges for: (1) CCfD strike prices and durations; (2) initial costs; and (3) learning-by-doing and learning-by-searching rates. Responses varied significantly, underscoring the lack of consensus around parameters that are fundamental to CDR learning. This uncertainty is central to the analysis, because different assumptions can lead to very different outcomes.

In order to prioritise our analysis, several policy sequences comprising timely policy instruments were compared based on the feedback from market participants in the questionnaire. The main comparison focused on a sequence which combined CCfD with later ETS integration, and a similar sequence that also includes a 20% research investment tax credit, as shown in Figure 4 below.

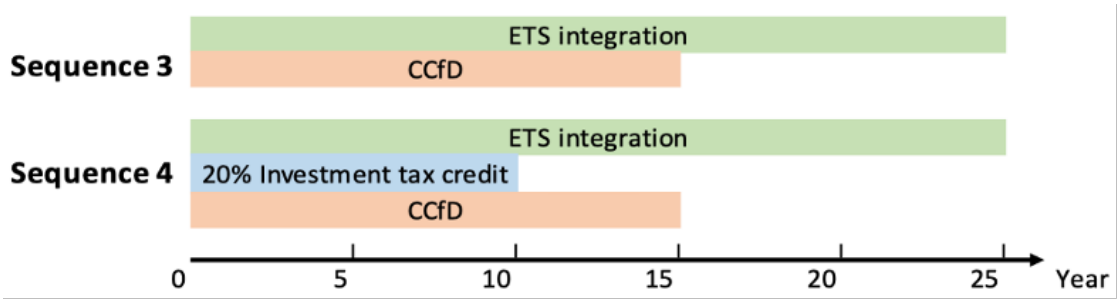


Figure 4: Policy sequences applied in the study based on market actors' feedback.

Using the range of elicited parameter values, the modelling results are presented as exploratory parallel-coordinate plots that make uncertainty and complexity transparent and illustrate the implication of endogenising learning within technology-learning simulations (see Figures 5 and 6). These plots allow analysts and stakeholders to compare many combinations of assumptions and policy settings at the same time, and to examine relationships between main input variables and deployment, cumulative removals and cost outcomes.

Applying the parallel plots, the modelling explores pessimistic, average and optimistic assumptions from elicitation inputs on initial costs, learning-by-doing and learning-by-searching rates, and policy support, as illustrated in Figure 5 below. The results suggest that:

- Strong early price support, especially through generous carbon contracts, can help drive deployment and cost reductions.
- Under pessimistic cost assumptions, very high policy support may be needed to induce meaningful removals.
- Under mid-range assumptions, deployment is possible but more limited. Adding an investment tax credit can strengthen R&D incentives and produce somewhat higher long-term removals.

Figure 5 is colour-coded according to the initial costs of DAC removal per tonne and shows how different learning rates, CCfD strike prices and contract durations impact DAC deployment in 2035 and 2050 as well as cumulative removals in 2050.

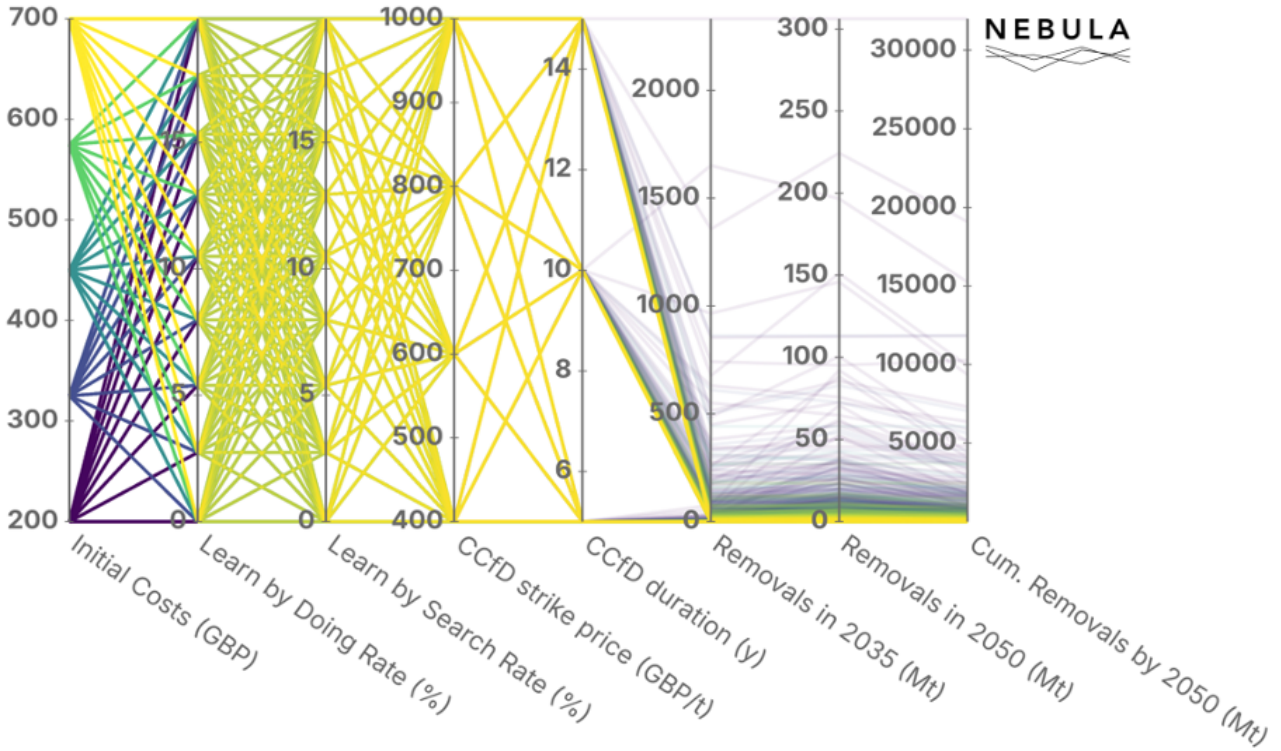


Figure 5: All possible scenarios under assumptions according to survey responses. **Note:** Input variables are initial costs, learning-by-doing, learning-by-searching, CCfD strike price and CCfD duration. Output variables are removals (per year) in 2035, removals (per year) in 2050 and cumulative removals by 2050. The numerical values are illustrative only and are intended to highlight the underlying dynamics, not specific values or magnitudes.

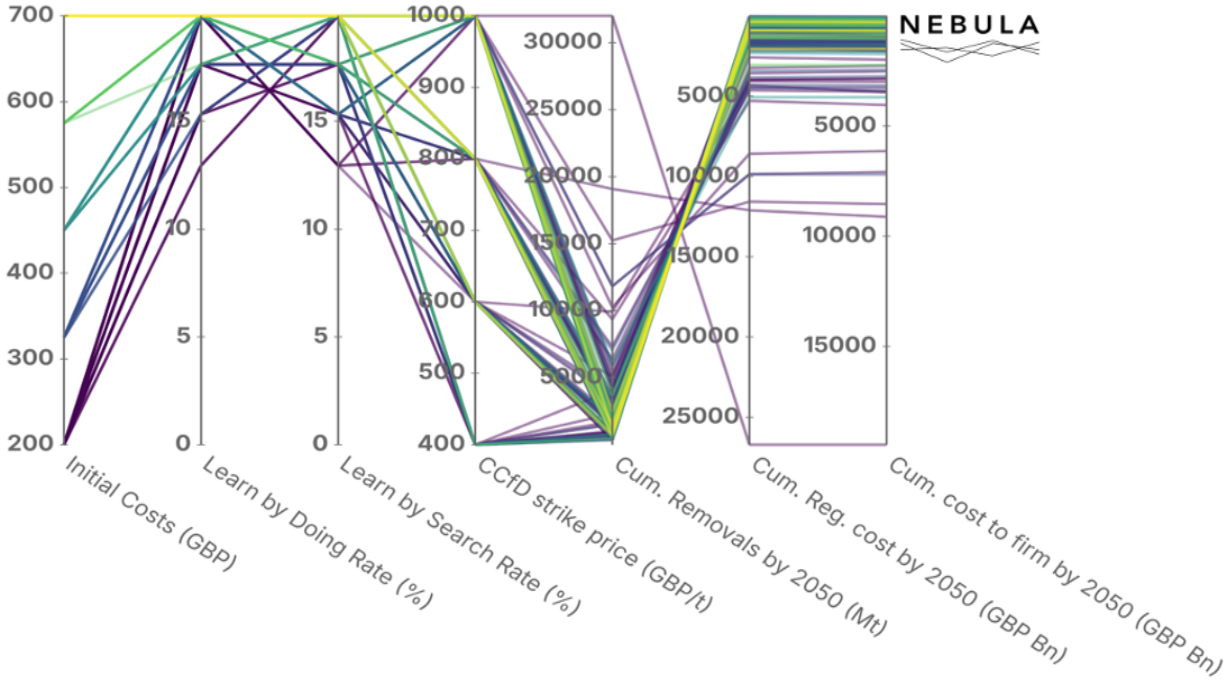


Figure 6: Scenarios achieve 5 MtCO2/year by 2035 and 10 MtCO2/year by 2050. **Note:** Input variables are initial costs, learning-by-doing, learning-by-searching and CCfD strike price. Output variables are cumulative removals by 2050, cumulative costs to regulators by 2050 and cumulative costs to the firm by 2050. The numerical values are illustrative only and are intended to highlight the underlying dynamics, not specific values or magnitudes.

The main takeaway is that learning can have a substantial effect on deployment and costs but does not eliminate the need for strong policy intervention. If CDR technologies can drive costs down through rapid learning, early support can accelerate movement towards lower costs and larger scaled deployment. However, if initial costs are high or learning rates are low, policy support may need to be much stronger and sustained for longer (i.e. a longer-duration CCfD facilitates more prolonged support, allowing further cost reductions). The large spread of possible outcomes cautions against relying on a single forecast or scenario, as is common in orthodox SFLC-based whole-systems analysis. The plots allow us to identify combinations and policy settings that achieve material DAC deployment by 2035 and 2050. In particular, we can examine scenarios consistent with illustrative deployment benchmarks set by the UK's Climate Change Committee. We therefore filter the scenarios shown in Figure 5 to include only those reaching 5 $MtCO_2$/year by 2035 and 10 $MtCO_2$/year by 2050, producing the subset shown in Figure 6. This allows stakeholders to compare scenarios that combine initial costs, learning-by-doing, learning-by-searching and different policy settings that achieve particular targets, such as 5 $MtCO_2$/year by 2035 and 10 $MtCO_2$/year by 2050.

4. Implications for UK DAC CDR development

After we had created the exploratory parallel computer visualisations produced from the model and made the uncertainty/complexity transparent across combinations of input variables, we then presented them at a workshop to ***engage a range of stakeholders in a deliberative set of exercises.*** The half-day workshop brought together 12 attendees from policy (DESNZ, LCCC); the third sector (Carbon Gap, Carbon Balance); DAC start-ups and developers (Mission Zero, Air Hive, OnePointFive, Progressive Energy); and finance (Seoi.Cure, HSBC, Green Finance Institute, Accountants for Sustainability).

The following perspectives emerged from the different groups:

The developer view: just get something built. For DAC developers, the most pressing problem is not the strike price of a future carbon contract; it is the sheer difficulty of getting any project to the point of final investment decision. The capital expenditure (CAPEX) requirement is enormous, funders are wary and many of the policy instruments currently on offer only kick in once a project is already operational and selling carbon credits.

Developers pointed to a specific timing problem: CCfDs may not be available until 2028 or the early 2030s, leaving an immediate gap for the first wave of projects. They called for CAPEX grants, state-backed performance guarantees and public procurement – mechanisms that reduce risk before a single tonne of CO_2 has been removed. One consistent message was that DAC will be financed through project finance, not corporate balance sheets, which makes the bankability of each individual project absolutely critical.

The finance view: the capital stack problem. Investors brought a different kind of clarity to the discussion. Working through a hypothetical £200 million first-of-a-kind (FOAK) DAC project, the finance group mapped out what a viable capital structure would actually look like. Their conclusion was stark: 50% of that sum would potentially need to come from government development funds or philanthropic capital willing to absorb first-loss risk. Only then would private credit funds consider the remaining portion. Traditional banks, with their near-zero tolerance for technical risk, would take only a small, heavily guaranteed slice of any investment.

The group also identified what an offtake agreement[33] needs to look like before any serious investor conversation can begin: secured, reputable, ideally collateralised and supported by secondary revenue streams. Without these foundations, no amount of policy generosity will make an investment feasible. Their broader contribution was that the government should act as initial insurer for FOAK projects. This is a role private markets could eventually take over once risks are better understood and quantified.

The policymaker and enabler's view: sequencing and fiscal reality. Policymakers and enablers brought the need for fiscal discipline and a longer strategic view to the decision space. They were broadly sympathetic to the case for early-stage support, but also clear about the fiscal limits of public budgets and political feasibility. On the central question of CCfDs, they agreed with developers and investors that locking in high strike prices now – before technology costs have fallen – risks not providing value for money for the taxpayer. Cost reductions via learning make this lock-in effect even more salient.

This group placed particular emphasis on demand-side instruments. ETS integration – tightening allowance caps to make DAC credits more competitive – was seen as essential, though the current UK ETS carbon price was widely acknowledged as too low to drive meaningful deployment. Mandates requiring a growing share of compliance obligations to be met through permanent removals were seen as a more durable way to create the demand signal that investors actually need, e.g. a sustainable aviation fuel mandate. This echoes the need to *'commit to additional support for pilot- and mid-scale projects to address "valley of death" issues and secure a pipeline of diverse GGRs'* spelled out in the Independent Review on Greenhouse Gas Removals.[34]

The following areas of convergence emerged across the groups:

- Policy sequencing matters as much as policy content. The right instrument at the wrong moment may achieve very little.
- Early deployment and de-risking must come before long-term market instruments such as CCfDs. CCfDs are better suited to scaling than to launching the first projects.

- Demand creation is the central policy challenge. Without credible, growing demand, no financing structure holds together.
- Government must act as first-loss investor and initial insurer for FOAK plants. The private sector cannot yet absorb that risk in the DAC sector.
- Technological diversity must be preserved. Policy should avoid inadvertently picking winners before the field has matured.

5. A deep uncertainty complexity approach to technology learning

None of these amounts to an easy answer for the establishment of DAC CDR technologies in the UK. The workshop surfaced deep tensions between fiscal prudence and the need for bold, early action; between revenue certainty for developers and value for money for taxpayers; and between technology-neutral policy and the practical need to make something happen soon.

In terms of the deep uncertainty complexity approach, the framework demonstrated that structured, deliberative dialogue between these groups produces something that the multiple factors involved in technology learning can stimulate when they are made transparent: a shared understanding of where the real obstacles lie and the complexity of the trade-offs that need to be considered when establishing a new sector from a standing start for which no precedent is available. Issues and dynamics that are assumed away in SFLC and whole-system optimisation modelling analysis stand to distort policy design for frontier technology development. This could not be more important, given the UK's increasingly constrained public finances and the need to demonstrate value for money.

Recommendations

- **Policymakers** need to be much more aware of the importance of the different aspects of technology learning when developing innovation policy, as well as the shortcomings of assumption-heavy frameworks that dominate the orthodoxy, which at best hide, and at worst distort, the relevant policy interventions required to fulfil innovation policy objectives.
- **Investors** should develop financing approaches that recognise the unique risk profile of early DAC projects, using blended finance and risk-sharing mechanisms to enable FOAK projects until technology and market risks are sufficiently reduced.
- **Technology developers** need to prioritise bankability alongside technical performance, recognising that learning-by-doing encompasses commercial, financial and operational capabilities as much as technological improvement.

Finally, coordination among policymakers, investors and technology developers will be essential to overcome the substantial barriers to widespread CDR deployment. The market is so nascent and the information asymmetries so substantive that neoclassical economic assumptions do not hold. There is a need for decision support frameworks that make the complexity transparent and facilitate the deliberation between parties to address these information asymmetries. The deep uncertainty approach we have put forward goes some way to addressing this. We therefore encourage academic institutions, consultancies and multilateral organisations involved in this space to further develop the deep uncertainty complexity approach to technology learning we have prototyped. We invite others to apply it, adapt it and critique it.

Acknowledgements

Insights for this article were based on research that was undertaken as part of the UK GGR-D programme and the Oxford University-led CO_2RE Hub: https://co2re.org/.

The authors would also like to thank Professor Cameron Hepburn of Oxford University and Dr Ajay Gambir of the Accelerator for Systemic Risk Assessment (ASRA) for reviewing the initial drafts of this policy brief.

Endnotes

1 Mark Workman is the primary and corresponding author: mark.workman07@imperial.ac.uk. Authors are listed in alphabetical order.

2 *Emissions Gap Report 2026*. UN Environment Programme. Published October 2026.

3 UK Climate Change Committee, 2025. *Seventh Carbon Budget*. Published February 2025.

4 See Direct Air Capture and Greenhouse Gas Removal Innovation Programme, 17 August 2020; and "New vision to create competitive carbon capture market follows unprecedented £20 billion investment". UK Government press release, 20 December 2023.

5 See Integrating greenhouse gas removals in the UK Emissions Trading Scheme, 23 May 2024; and GGRs and Power BECCS business models: amending revenue support regulations relating to directions to a counterparty, publication of information, counterparty's ability to carry out its functions, and eligibility, April 2025.

6 Bistričić, P., et al. 2025. *Unlocking Private Investment for Greenhouse Gas Removal: Public-private collaboration and a technology-specific policy approach*. CO_2RE policy briefing. Published September 2025.

7 Ragazzi, I., et al. 2025. *Update on Greenhouse Gas Removal Costs and Scaling Challenges*. Joint Report with ERM and CO_2RE. Initial ERM analysis in 2021 did not accommodate for path dependence of technology development, i.e. that not all policy instruments will have an equal effect at a single point in time and therefore modelling implies that they are inter-temporally discrete.

8 Workman, et al. 2026. "Investing in carbon dioxide removals: A new analytical and policy paradigm". *Energy Research & Social Science* 135 (2026) 104669.

9 Indeed, the CDR sector is considered to be steeped in the characteristics of deep uncertainty. This is a situation *where analysts do not know, or the parties to a decision cannot agree on, (1) the appropriate conceptual models that describe the relationships* among *the key driving forces that will shape the long-term future, (2) the probability distributions used to represent uncertainty about key variables and parameters in the mathematical representations of these conceptual models, and/or (3) how to value the desirability of alternative outcomes. In particular, the long-term future may be dominated by factors that are very different from the current drivers and hard to imagine based on today's experiences – see endnote 6.*

10 Beinhocker, E., and Bednar, J. 2025. *Complexity and Paradigm Change in Economics*. INET Oxford Working Paper Series, No. 2025-20 – see page 485.

11 Workman, M., et al., 2021. "Climate policy decision making in contexts of deep uncertainty – from optimisation to robustness". *Environmental Science & Policy* 120 (2021) 127–137.

12 Consolidative modelling is where analysts consolidate known facts to generate a detailed surrogate of the real world. This contrasts with exploratory modelling, which samples over ensembles of computational experiments that represent plausible assumptions about the future and how it may unfold.

13 Edwards, M., et al. 2026. *The State of Carbon Dioxide Removal 3rd Edition*.

14 Trancik, J.E., et al. 2025. "Informed investments in clean energy technologies". *Nature Energy* Volume 10 December 2025, 1404–1411.

15 Feng, S, Stemmler, J. and Workman, M. (forthcoming). "Policy Sequencing for Emerging Technologies: Demand Creation and Endogenous Learning".

16 Workman, M., et al., (forthcoming). "Multi-factor Technology Learning Curves – a deep uncertainty approach to policy design". To be submitted to *Environmental Science & Policy*.

17 Grubler A., and Wilson, C., 2013. *Energy technology innovation: learning from historical successes and failures*. Cambridge: Cambridge University Press.

18 Farmer, J.D. and Lafond, F., 2016. "How predictable is technological progress?" *Research Policy* 45 (2016) 647–665.

19 Lafond F., 2025. "Forecasting technological progress". In: del Rio-Chanona, R.M., et al. (eds). *The Economy as an Evolving Complex System IV*. The Santa Fe Institute Press, 2025.

20 Way, R., et al., 2022. "Empirically grounded technology forecasts and the energy transition". *Joule* 6, 2057–2082.

21 *The Economist*, 26 March 2022. "Britain's next nuclear plant will be identical to one under construction".

22 UKERC, 2015. "Innovation timelines from invention to maturity".

23 Nemet, G.F., et al. 2023. "Near-term deployment of novel carbon removal to facilitate longer-term deployment". *Joule*, vol. 7, no. 12, pp. 2653–2659.

24 Meissner, L. and Burke, J. 18 June 2026. "Will carbon removal ever be as cheap as solar?" Grantham Institute for Climate Change and the Environment Commentary.

25 Two Integrated Assessment Models do apply two factor learning curves in their analysis. These are WITCH and REMIND-MAgPIE, which break out innovation (R&D and knowledge spillovers) and diffusion via learning-by-doing. These are applied to mitigation technologies and the extent to which they are endogenously applied is unclear. For example, REMIND-Mag-PIE states that for a basket of mitigation technologies specific investment costs decrease by 12, 20, and 9%, respectively, for each doubling of cumulated capacity, and capital costs decrease with a 10% learning rate for others. The application of two-factor learning and the ambiguity as to how they are applied aside, the inherent lack of transparency as to how technology learning is applied and their impacts on technology deployment in orthodox net-zero analysis remains.

26 Gambhir, A., et al. 2021. "How Are Future Energy Technology Costs Estimated? Can We Do Better?" International Review of Environmental & Resource Economics, 2021, 15: 1–48.

27 Lafond, F., Greenwald, D., and Farmer, J.D. 2022. "Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment". *The Journal of Economic History*. 2022;82(3):727-764.

28 Elia, A., et al. 2021. "Impacts of innovation on renewable energy technology cost reductions". *Renewable and Sustainable Energy Reviews* 138 (2021) 110488.

29 Learning-by-doing relates to cost reductions achieved due to experience gained during the production process, such as the improvements in management and labour knowledge acquired. Learning-by-using relates to improvements generated through feedback from users during operational activities. Learning-by-interaction refers to cost reductions achieved through knowledge exchange with other actors along the technology value chain, including benchmarking between firms, industries and research centres. It may also include knowledge exchanged between locations, for example through joint-venture collaborations in new markets. Knowledge spillover may also arise from advances in related technologies, which can be understood as the inter-industry spillover. Factors such as learning-by-competition and x-efficiency, which arise from competitive pressure, managerial discipline or market structure rather than direct knowledge exchange, should be treated as market-related drivers.

30 Yu, C.F., van Sark, W.G.J.H.M., Alsema, E.A. 2010. "Unraveling the photovoltaic technology learning curve by incorporation of input price changes and scale effects". *Renewable and Sustainable Energy Reviews* 2011; 15:324–37.

31 Nemet, G.F. 2006. "Beyond the learning curve: factors influencing cost reductions in photovoltaics". *Energy Policy* 2006; 34:3218–32.

32 City Science and Imperial College London Consultants. 2025. *Assessing the Feasibility for Large-scale DACCS Deployment in the UK*. Report for UK Climate Change Committee.

33 An offtake agreement is a legally binding contract between a producer and a buyer to purchase goods yet to be produced. These agreements are crucial in project financing as they assure future income, making it easier to secure loans for new facilities or expansions. Common in volatile markets, offtake agreements lock in prices and ensure supply for buyers, providing financial stability for producers.

34 Whitehead, A. 2025. *Independent Review of Greenhouse Gas Removals*. Department for Energy Security and Net Zero.

CO$_2$RE THE GREENHOUSE GAS REMOVAL HUB

www.co2re.org

Funded through the UKRI SPF Greenhouse Gas Removal Demonstrators (GGR-D) Programme, CO$_2$RE co-ordinates the Programme and conducts solutions-led research to evaluate a balanced portfolio of economically, socially and environmentally scalable Greenhouse Gas Removal options, with associated policy design, engagement and outreach.

UK Research and Innovation

CO$_2$RE , The Greenhouse Gas Removal Hub is funded by UK Research and Innovation grant reference NE/V013106/1.