The Carbon at Risk measure can unlock financial markets for gigaton-scale carbon removal

Bali Lee[1,2], Nick Gogerty[3], Oleksandr Kit[3], Delia Meth-Cohn[1], Frederic Olbert[4], Frank Venmans[5], Gabrielle Walker[1], Paul Young[6], and Ben Groom[5,7]

[1]Rethinking Removals, UK.
[2]Climate Resilient Solutions, UK.
[3]Carbon Finance Labs, US.
[4]CarbonPool, Switzerland.
[5]Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, UK.
[6]Kita, UK.
[7]Dragon Capital Chair in Biodiversity Economics, LEEP Institute, Department of Economics, University of Exeter Business School, Exeter, EX4 4PU, UK. Email: b.d.groom@exeter.ac.uk. Corresponding author.

ABSTRACT

This is the Supplementary Information for the paper "The Carbon at Risk measure can unlock financial markets for gigaton-scale carbon removal".

Supplementary Information

S1. Graphic Definition of Carbon at Risk (CaR)

Carbon at Risk definition

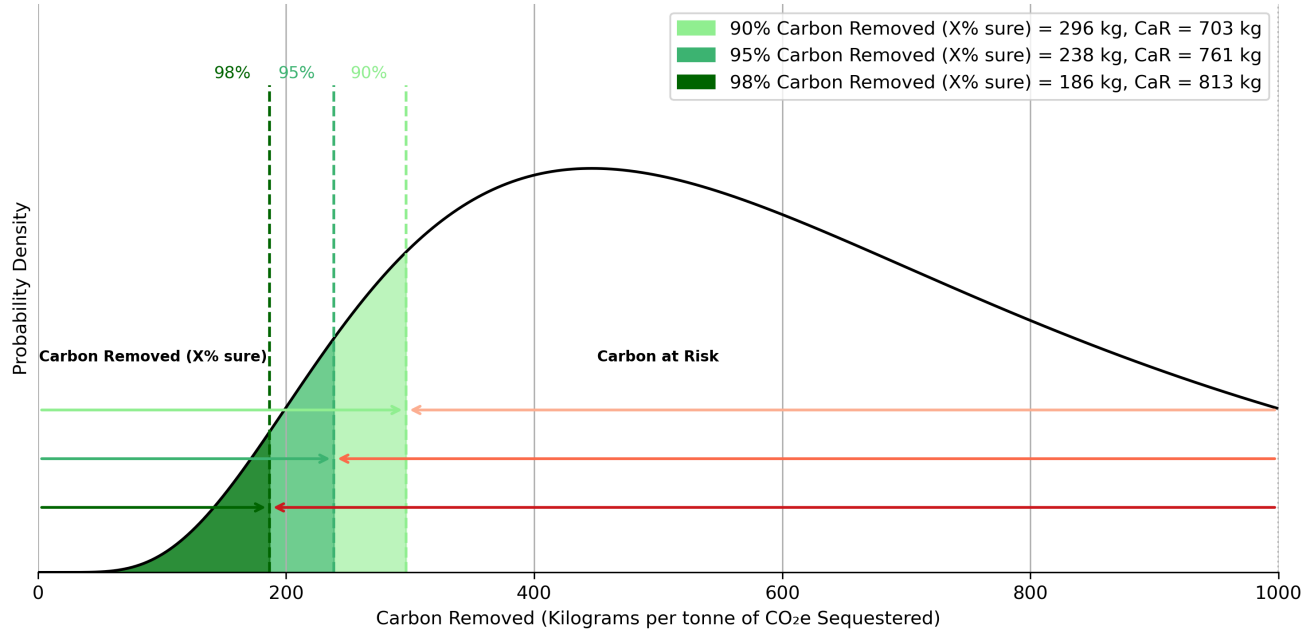

Figure S1. The Carbon at Risk and Carbon removed (X% sure): A less reliable carbon removals technology than in Figure 1b in the main text.

S2. Details on the Methodology of CaR Assessment for Nature-Based Carbon Sequestration in Forests: Data Processing Steps and Decisions

Assumptions and Parameters

The basic assumptions and parameters surrounding estimation of the CaR for forests in the main text are as follows:

Parameter	Value
Project area	10,000 hectares
Annual climate change-induced burn rate increase	0.5%
Time horizon of the analysis	100 years
Number of Monte Carlo simulations	1,000
Upper Gross Primary Productivity regrowth threshold	2,500 g C/m^2/year

Table S1. Parameters for the forest CaR simulation

The Jupyter Notebook for Physical Risks and Regrowth Calculations

A Jupyter Notebook was created to allow simulations for regions across the globe. This link to the Notebook is available on request from the authors. The Notebook is initialized by selecting up to three countries and sub-national administrative units for the analysis as in Figure S1 below.

Then, for the selected administrative units, the annual burned area and forest cover data can be accessed from the EU Global Wildfire Information System (GWIS) for years 2002 through 2023 through the GWIS API. Then, for each administrative unit the average Gross Primary Productivity (GPP) is calculated using data from[1]. Figure S2 shows the resulting maps of GPP. GPP is used as the principal factor for post-fire regrowth as per[2], as shown in Figure S3.

Regional land cover and fire history for the selected locations is displayed as in Figure X in the main text, reproduced here as Figure S4 for convenience.

Details of Fire Risk Carbon Loss and Regrowth Calculations

A project area adjustment coefficient is used to normalize across the simulations using a 1..3 ratio between project size and the average fire size (firesize field from GWIS API). This ensures that the risk to projects that are much larger than the average fire event size in the region is lower than the risk for a project that can be fully covered by a single fire event.

The burn area fraction is computed as a relationship between wildfires that occur in forested areas and the total forest area of the administrative unit using the `lc1` field from the GWIS API, multiplied by the project area adjustment coefficient.

Finally, Monte Carlo simulations are performed to obtain the 95th percentile of burned area fractions for individual administrative units over the course of 100 years, applying the annual regrowth factor every year. When plotted this results in figure X in the main text, reproduced here as Figure S5 for convenience.

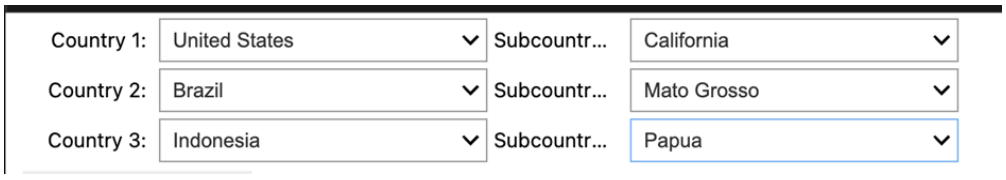

Figure S2. Initialization of the Jupyter Notebook

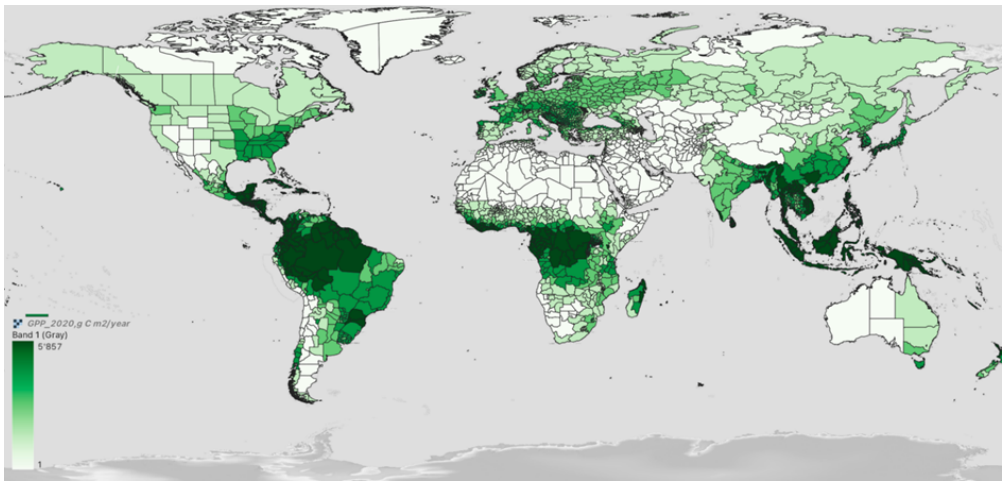

Figure S3. Spatial distribution of the annual average GPP from 1992 to 2020: Data source[1]

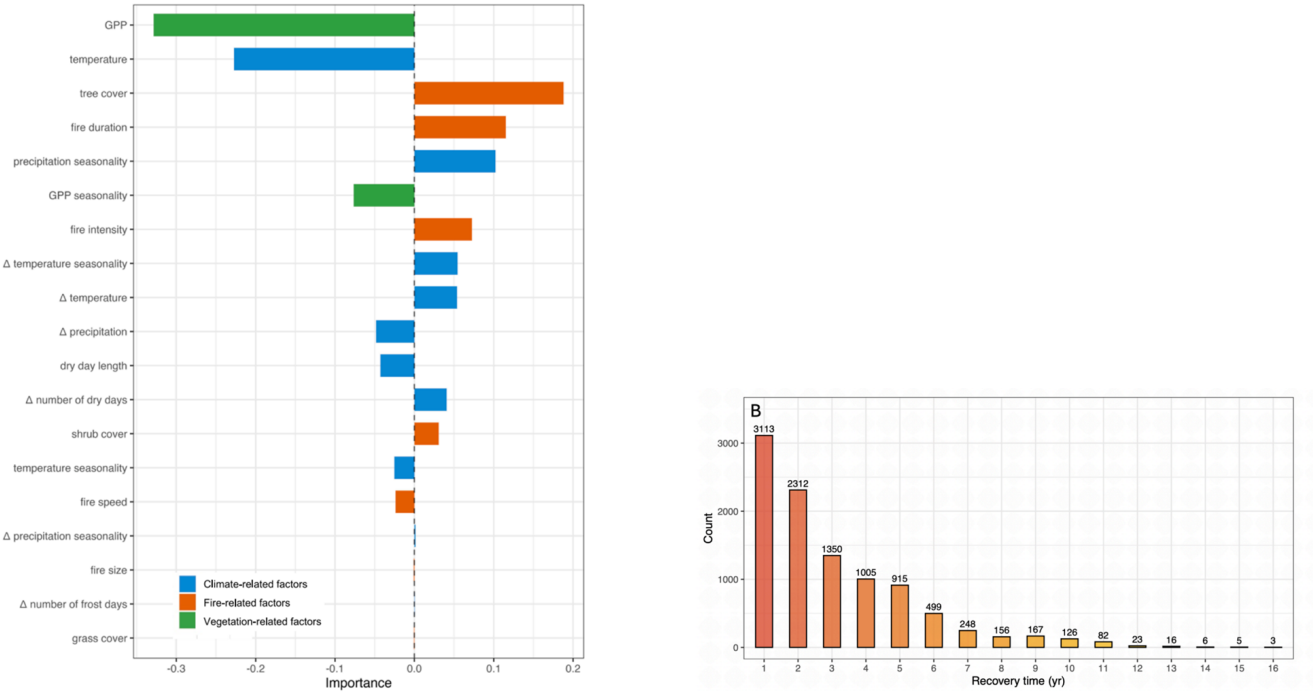

(a) Importance Factors **(b)** Recovery Time

Figure S4. Importance of factors and recovery times impacting post fire recovery: Source:[2]

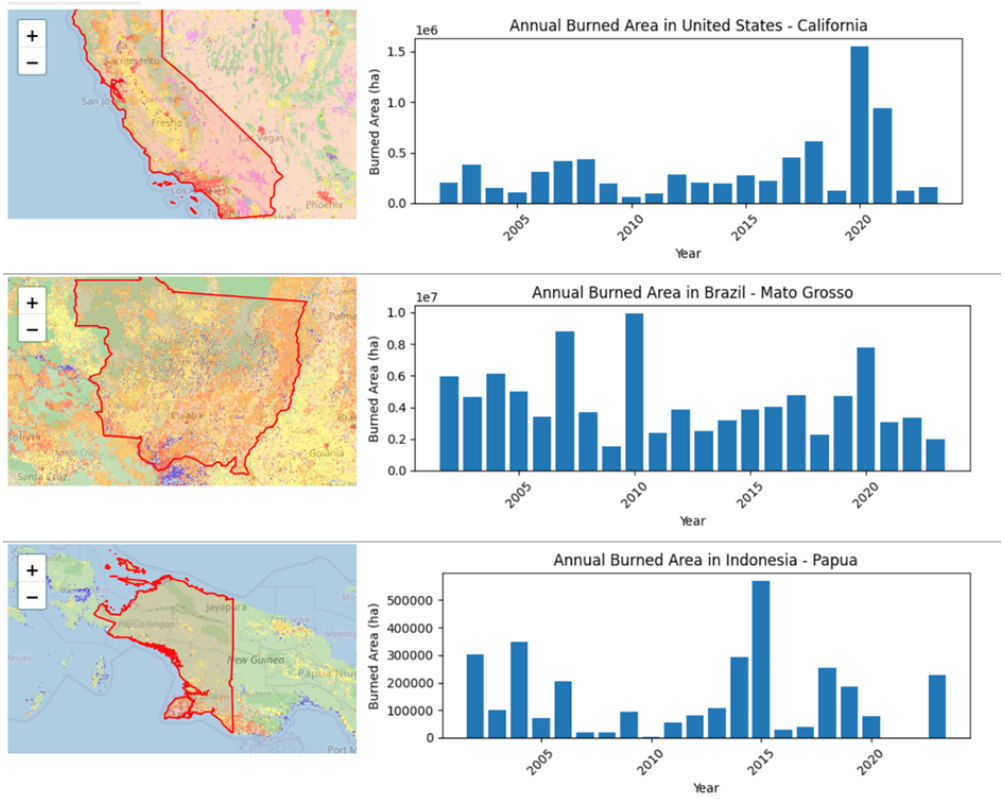

Figure S5. Regional maps and wildfire patterns: source: GWIS

Project	**Per tonne (kg/t C)**			**Total (tC)**		
	1 yr	10 yrs	100 yrs	1 yr	10 yrs	100 yrs
United States - California	191	524	788	63,720	174,808	262,788
Brazil - Mato Grosso	131	335	504	43,561	111,647	167,937
Indonesia - Papua	17	28	47	5,623	9,257	15,606
Portfolio Total	—	—	—	**112,904**	**295,712**	**446,331**

Table S2. Carbon at Risk: Per-ton and Total estimates for selected projects.

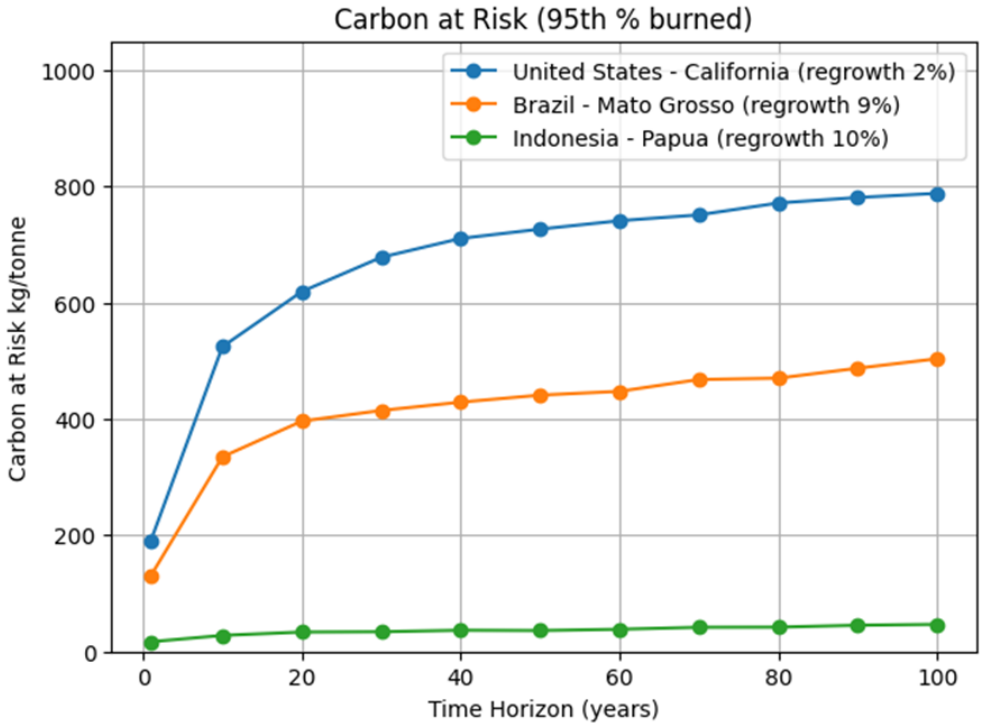

Figure S6. 95% Carbon-at-Risk curve by Region

Caveats to the Forest Carbon Analysis

Fires reported by GWIS use the MODIS burned area product (MCD64A1) to define wildfire events and compute the burnt area of each event. Apart from genuine wildfires this includes agricultural fires (slash-and-burn), controlled burns, and other open fire events that are not forest fires. This overestimates fire hazard considerably. This problem can be mitigated by applying spatio-temporal clustering algorithms, which are not part of this study. This problem is currently partially addressed by using the average fire size field from GWIS in the project area adjustment coefficient, as agricultural slash-and-burn fires tend to be much smaller than the "real" wildfires. For simplification reasons, this approach does not take into account project management risks or political and socio-economic landscape. We provide the basic starting point for more complicated analysis in the future.

Credits and CaR from avoided methane leakage

As a supplementary illustration of the way in which CaR can be used, in this section we articulate an example that looks at the risks of methane leakage evidence from different reticulation networks. In the context of DACCS this example could be used as a rough proxy to estimate the risk associated with delivering CO_2 to geological or other storage (delivery risk), noting that there are reasons why leakage of CO_2 and methane might be higher or lower from any given network, due to higher operating pressures and molecule size[3].

We use publicly available data on methane leakage from distribution networks as a proxy estimate of the probability density function (likelihood) of CO_2 losses in the implementation of DACCS[4,5]. Using the raw data on the frequency of leakage events of different sizes reported in studies[4,5] we estimate the parameters of a log normal distribution for leakage, which is a good approximation for events of this nature. By identifying the entire probability distribution in this way we are able to calculate the CaR for any confidence level. We then obtain the delivery risk CaR over time (as in Figure 2a) using Monte Carlo analysis, modeling each vintage of CO_2 as facing discrete transport risk during its delivery phase before safe storage. The results are shown in Fig S7 for two separate networks of different vintages.

Fig S7a shows the empirical distributions of leakage events for an existing network (in blue) and a newer network with a greater use of plastic piping (in green). The latter technology clearly reduces the risk of large leakage events, illustrated by the reduced spread of the distribution and the lower probability mass associated with large leakage events. Furthermore, in terms of the 95% CaR, these two networks are very different. The newer network reduces the the 95% CaR from 21% of the total distributed gas to approximately 11%. In units of CO_2e this translates as a drop from losses from 70.6MtCO_2e to 34.6MtCO_2e, a drop of 51%. Figure S7b takes the parametric estimates of the log-normal approximation of the distributions in Fig. S7a and

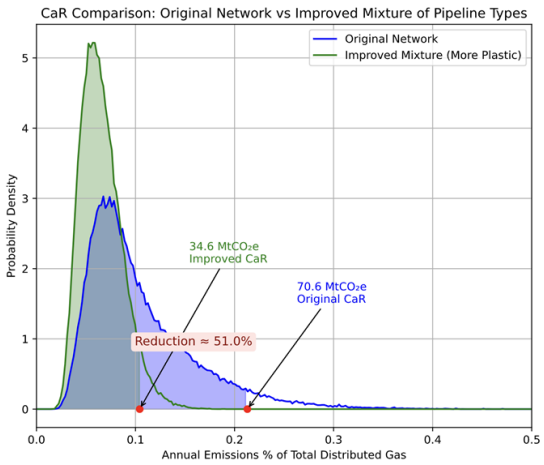

(a) 95% CaR methane network leakage.

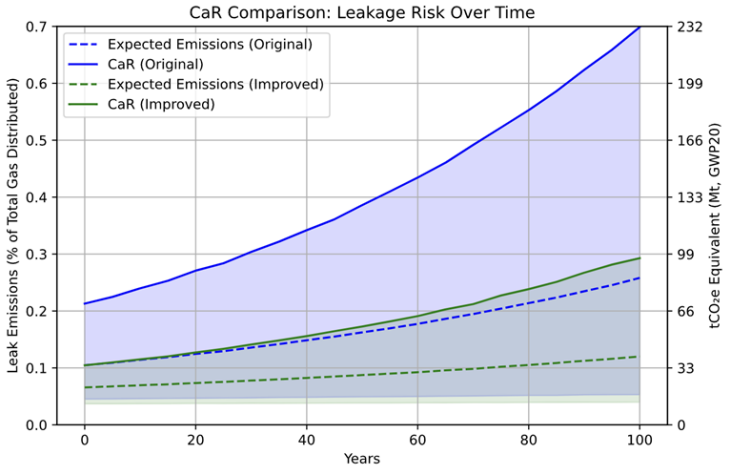

(b) CaR over time for methane network leakage.

Figure S7. 95% Carbon at Risk in (equivalent) tonnes per kilogram for methane network leakage: (a) the probability distribution for two methane transport networks, one old (blue) one newer (green). The CaR (Carbon equivalants) are 70.6 $MtCO_2e$ and 34.6 $MtCO_2e$ respectively. The approach taken uses data from[4] and[5]. (b) shows how leakage accumulates over time in each case, and how the relative value of leakage control and newer transport systems increases over time: the gap between the CaR and the expected values for each network increases over time.

projects the cumulative 95% CaR over a 30-year operational horizon. The CaR increases as new CO_2e vintages are processed through the networks each year of facility operation. The gap in CaR between new and older transport technology (in this case pipes) increases over the operational lifetime, reflecting the cumulative benefit of reduced network losses across all processed vintages. A similar gap emerges when comparing the expected (mean) values (the dotted lines in Fig S7b), illustrating the dynamic gains in risk reduction obtained from investment in reducing network leakage

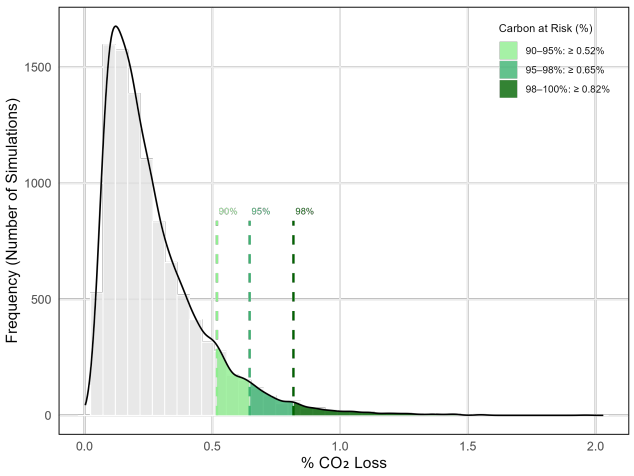

(a) Offshore DACCS 95% CaR for 1000yr horizon.

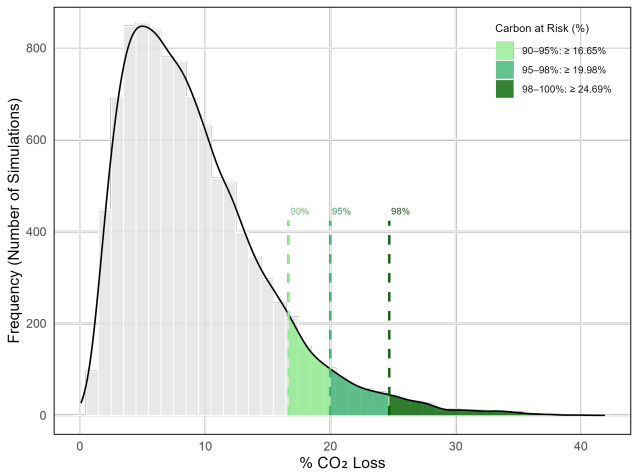

(b) Onshore, unregulated DACCS 95% CaR for 1000yr horizon.

Figure S8. Carbon at Risk for Direct Air Capture and Carbon Storage technologies. We use the Storage Security Calculator Monte Carlo analysis code from[3] to simulate the distribution of leakage as a percentage of overall storage for a 1000 year time horizon. The 90, 95 and 98% CaR are reported for 10000 repetitions of their "well-regulated offshore" scenario (a), the scenario with the most integrity, and the "poorly regulated onshore" scenario (b), the scenario with the least integrity. Note the differences in the scale of the X-axis.

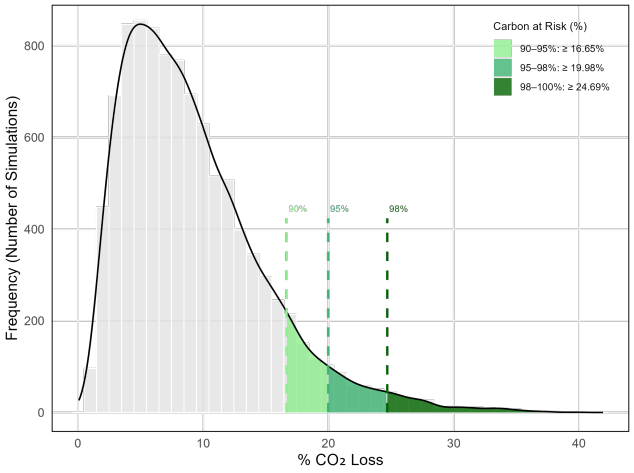

(a) Offshore DACCS 95% CaR for 10000yr horizon.

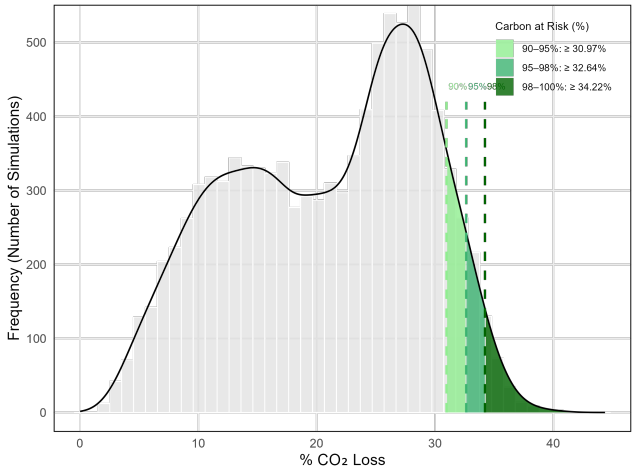

(b) Onshore, unregulated DACCS 95% CaR for 10000yr horizon.

Figure S9. Carbon at Risk for Direct Air Capture and Carbon Storage technologies. We use the Storage Security Calculator Monte Carlo analysis code from[3] to simulate the distribution of leakage as a percentage of overall storage for a 10000 year time horizon. The 90, 95 and 98% CaR are reported for their "well-regulated offshore" scenario (a), the scenario with the most integrity, and the "unregulated onshore" scenario (b), the scenario with the least integrity.

Portfolio simulations: Risk - return, diversification and costs

The cost per ton of independent portfolios and the p_5/CaR constraint.

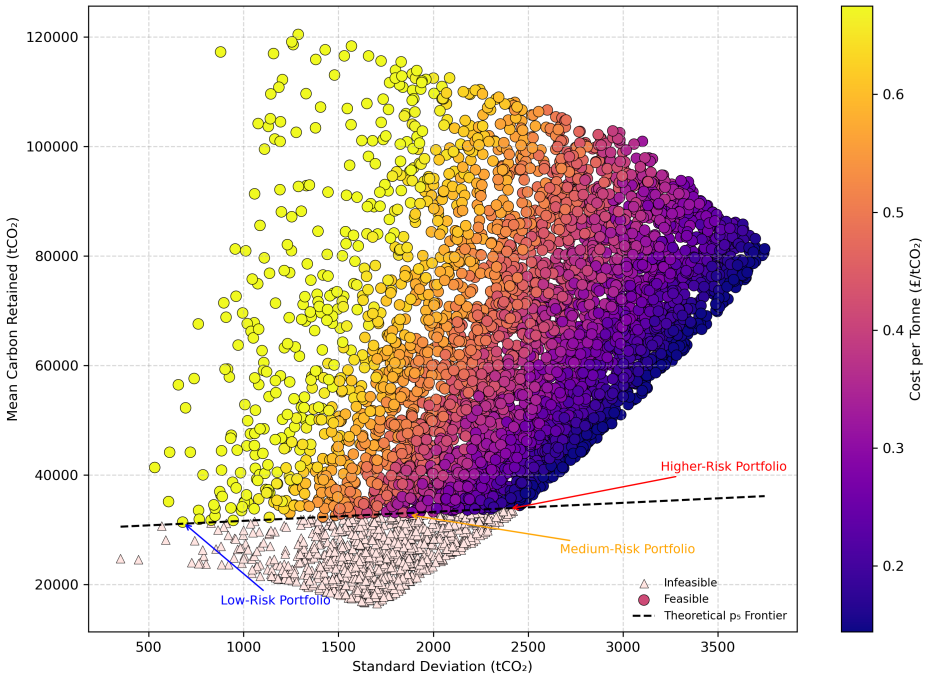

Figure S10. Cost per ton of Portfolios: In this simulated example the portfolios are made up of the three different technologies: 1) Low CaR DACCS (1% loss after 200 years, \$400/tCO$_2$); 2) Medium CaR BioChar (20% loss after 200 years, \$120/tCO$_2$); 3) High CaR forest carbon (35% loss after 200 years, \$40/tCO$_2$). The risk (standard deviation) and return (mean) trade-off as portfolios change in their project composition and the sets of portfolios that either do (circle), do not (triangle) or just meet the p5 requirement of 95% durable carbon after 200 years of 30000 tonnes (red circles). The shading shows the cost per unit of carbon removed associated with each portfolio. In this simulation portfolios with more forest carbon are cheaper despite the higher risk.

Risk–return framing of carbon removal projects

We model carbon removal portfolios within the mean–variance framework of Markowitz[6,7]. For each technology type $i \in \{1,\dots,N\}$ we specify:

- cost per tonne c_i (\$/tCO$_2$),
- survival probability π_i (expected permanence),
- survival variance σ_i^2 (variance of the delivery fraction),
- project size q_i (tCO$_2$ per project).

For budget B, the tonnes contracted from type i are

$$Q_i = \frac{B}{c_i}, \qquad n_i = \frac{Q_i}{q_i} \quad \text{projects.}$$

The expected delivered tonnes are

$$\mu_i = Q_i \pi_i = \frac{B}{c_i}\,\pi_i.$$

Variance–covariance structure

We assume outcomes are jointly normal with two layers of correlation:

- **Within–type correlation** ρ_i^{within} between the n_i projects of technology i,
- **Between–type correlation** $\rho_{ij}^{\text{between}}$ across technologies $i \neq j$.

Within–type variances. For technology i,

$$\text{Var}(X_i) \;=\; \left(\frac{B}{c_i}\right)^2 \frac{\sigma_i^2}{n_i}\Big(1+(n_i-1)\rho_i^{\text{within}}\Big).$$

Between–type covariances. For $i \neq j$,

$$\text{Cov}(X_i,X_j) \;=\; \frac{B}{c_i}\frac{B}{c_j}\,\sigma_i\sigma_j\,\rho_{ij}^{\text{between}}.$$

Matrix form. Stacking $X=(X_1,\ldots,X_N)$, we write

$$X \sim \mathcal{N}(\mu,\Sigma), \qquad \Sigma_{ii}=\text{Var}(X_i), \quad \Sigma_{ij}=\text{Cov}(X_i,X_j).$$

Correlation structure: modelling choices

A natural first attempt was a "fully negative" scenario, where all pairwise correlations are set to $\rho_{ij}^{\text{between}} = -0.6$ for $i \neq j$:

$$\rho^{\text{between}} = \begin{bmatrix} 1 & -0.6 & -0.6 \\ -0.6 & 1 & -0.6 \\ -0.6 & -0.6 & 1 \end{bmatrix}.$$

This assumption exaggerates diversification but produces a variance– covariance matrix Σ that is nearly singular. The efficient frontier then contains flat faces and sharp corners, and numerical optimisation (SLSQP) tends to cut inside the feasible set, so some Monte Carlo portfolios appear above the frontier. It is also implausible that all technologies perfectly hedge one another.

A more realistic specification is to allow only *forests* to be negatively correlated with DACCS and biochar, while keeping DACCS and biochar uncorrelated:

$$\rho^{\text{between}} = \begin{bmatrix} 1 & 0 & -0.6 \\ 0 & 1 & -0.6 \\ -0.6 & -0.6 & 1 \end{bmatrix}.$$

In this case forests act as the volatile hedge asset, while DACCS and biochar form a stable core that do not offset one another. The variance– covariance matrix remains well conditioned, the efficient frontier bends outward due to the hedge from forests, and the theoretical frontier aligns closely with the Monte Carlo cloud. This is also economically plausible: forest permanence may be reduced by climatic shocks in years when DACCS powered by renewables performs well, while DACCS and biochar share energy and supply chain risks. We therefore adopt this structure in the main analysis, presenting the fully negative case only as a limiting thought experiment.

Portfolio variance expansion

Let $w=(w_D,w_B,w_F)$ be the portfolio weights for DACCS (D), Biochar (B), and Forests (F). The portfolio outcome is

$$X_p = w_D X_D + w_B X_B + w_F X_F.$$

The variance is

$$\text{Var}(X_p) = w^\top \Sigma w,$$

which expands to

$$\begin{aligned}\text{Var}(X_p) = {} & w_D^2\,\text{Var}(X_D) + w_B^2\,\text{Var}(X_B) + w_F^2\,\text{Var}(X_F) \\ & + 2w_Dw_B\,\text{Cov}(X_D,X_B) + 2w_Dw_F\,\text{Cov}(X_D,X_F) + 2w_Bw_F\,\text{Cov}(X_B,X_F).\end{aligned}$$

Here:

- The squared terms reflect own–variances after within–type scaling.
- The cross terms reflect between–type correlations.

Monte Carlo simulation

To visualise the feasible set, we simulate random portfolios by drawing weights w uniformly on the simplex ($w_i \geq 0, \sum_i w_i = 1$). For each portfolio, we compute

$$R = \mu^\top w, \qquad \sigma_p = \sqrt{w^\top \Sigma w}.$$

Within–type assumptions in Monte Carlo. Each technology's variance term $\mathrm{Var}(X_i)$ already incorporates the number of projects n_i and its within–type correlation ρ_i^{within}. For example, if forests deliver in 500 t blocks, then $n_F = Q_F/500$ and the diagonal term for forests is reduced accordingly (unless $\rho_F^{\text{within}} = 1$, in which case there is no reduction). Thus the simulation respects the project size assumption.

Between–type assumptions in Monte Carlo. The off–diagonal elements of Σ encode the chosen correlation scenario ($\rho_{ij}^{\text{between}}$). Thus each draw of w produces a risk estimate that reflects whether technologies are positively, independently, or negatively correlated.

Interpretation. The Monte Carlo cloud is therefore not "just random" but already contains both the within– and between–correlation assumptions. The efficient frontier appears as the upper envelope of the cloud.

Parameterisation (example)

For simulation we set $B = \$1\text{m}$. Illustrative parameters:

$$\begin{aligned}
&\text{DACCS:} && c = 400\ /\text{t},\ \pi = 0.99,\ \sigma = 0.02,\ \rho_{\text{D}}^{\text{within}} = 1;\\
&\text{Biochar:} && c = 120\ /\text{t},\ \pi = 0.80,\ \sigma = 0.10,\ \rho_{\text{B}}^{\text{within}} = 0.5;\\
&\text{Forests:} && c = 40\ /\text{t},\ \pi = 0.65,\ \sigma = 0.25,\ \rho_{\text{F}}^{\text{within}} = 0.3.
\end{aligned}$$

Between–type correlation scenarios use $\rho_{ij}^{\text{between}} \in \{+0.6, 0, -0.6\}$ to illustrate positive, zero, and negative co-movement. In the negative case we adopt the "forests–negative" specification described above, with forests negatively correlated to DACCS and biochar, while DACCS and biochar remain uncorrelated. These choices yield expected removals of 2,475; 6,667; and 16,250 tCO_2 for DACCS, biochar, and forests respectively. Diagonal entries of Σ reflect within–type scaling via n_i and ρ_i^{within}, while off–diagonals reflect the chosen $\rho_{ij}^{\text{between}}$.

Summary

This framework translates heterogeneous features of carbon removal technologies — costs, survival probabilities, delivery variances, project granularity, and correlations — into a unified risk–return model.

Two distinct correlation layers are made explicit:

- *Within–type correlation* ρ_i^{within} determines whether multiple projects of the same technology behave like many independent draws (variance scales as $1/n_i$) or a single perfectly correlated block (no variance reduction). Project size q_i sets the number of projects n_i under a given budget B, which together with ρ_i^{within} fixes the diagonal entries of Σ.
- *Between–type correlation* $\rho_{ij}^{\text{between}}$ determines how technologies hedge one another. Positive values compress the frontier inward, zero values produce the classical $\sigma/\sqrt{n}$ diversification law, and negative values shift the frontier outward through hedging benefits. These correlations set the off–diagonal entries of Σ. In practice we adopt the "forests–negative" specification as the benchmark negative case, since it is both numerically stable and economically plausible.

Given expected returns μ and covariance matrix Σ, the Markowitz programme identifies efficient portfolios that minimise delivery risk for each expected removal target. Monte Carlo simulations visualise the feasible set, showing how correlation assumptions shape the opportunity set of carbon removal portfolios.

In short, the method demonstrates how cheap but volatile technologies (forests) can be mixed with expensive but robust ones (DACCS) to manage long–term carbon delivery risk. The role of correlation is central: within–type correlation controls the extent of variance reduction inside a technology, while between–type correlation governs hedging possibilities across technologies.

Efficient mean-risk frontier under different correlation scenarios

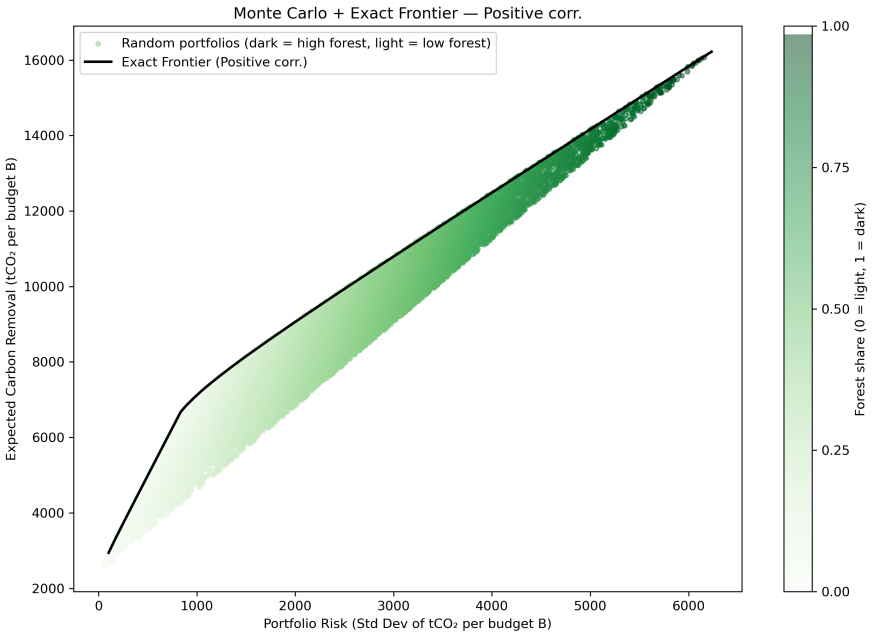

Figure S11. Monte Carlo cloud of random portfolios (shaded by forest share) and the exact Markowitz efficient frontier (black) under positive correlation of project risks. Diversification benefits are limited and forests only enter at the high–risk end.

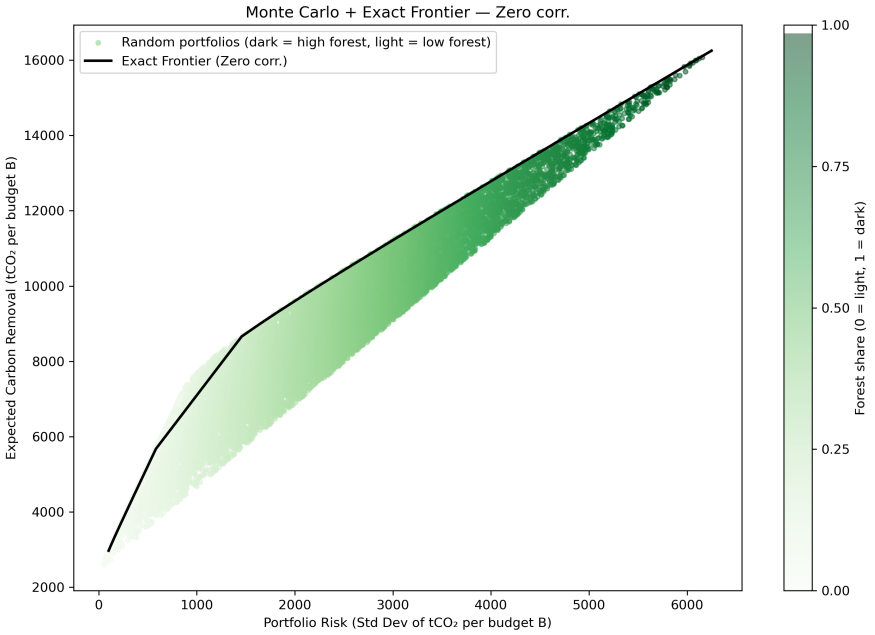

Figure S12. Monte Carlo portfolios and efficient frontier under zero correlation. Diversification reduces risk substantially, with biochar entering portfolios at intermediate risk levels.

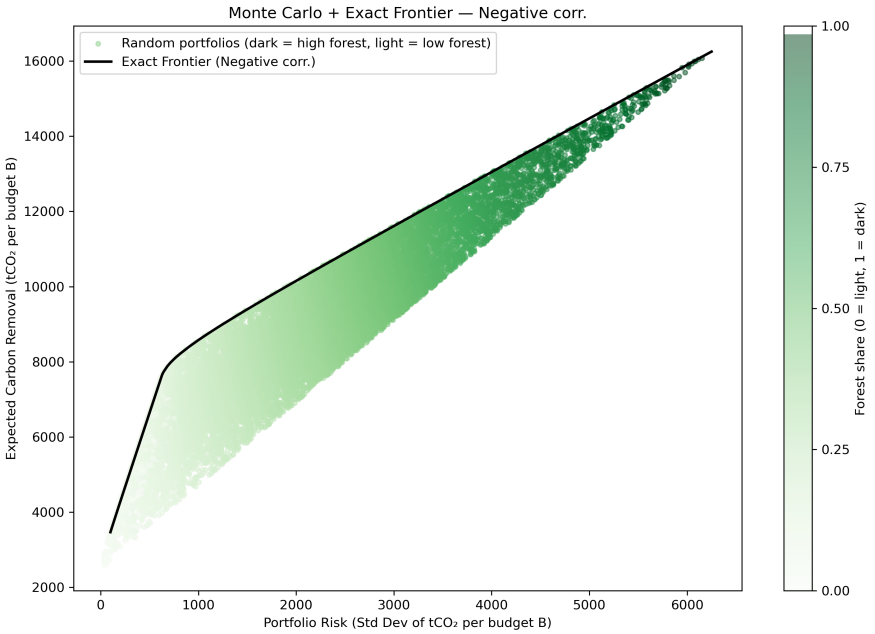

Figure S13. Monte Carlo portfolios and efficient frontier under negative correlation. Forest projects are included even at moderate risk due to hedging effects, pushing the frontier outward.

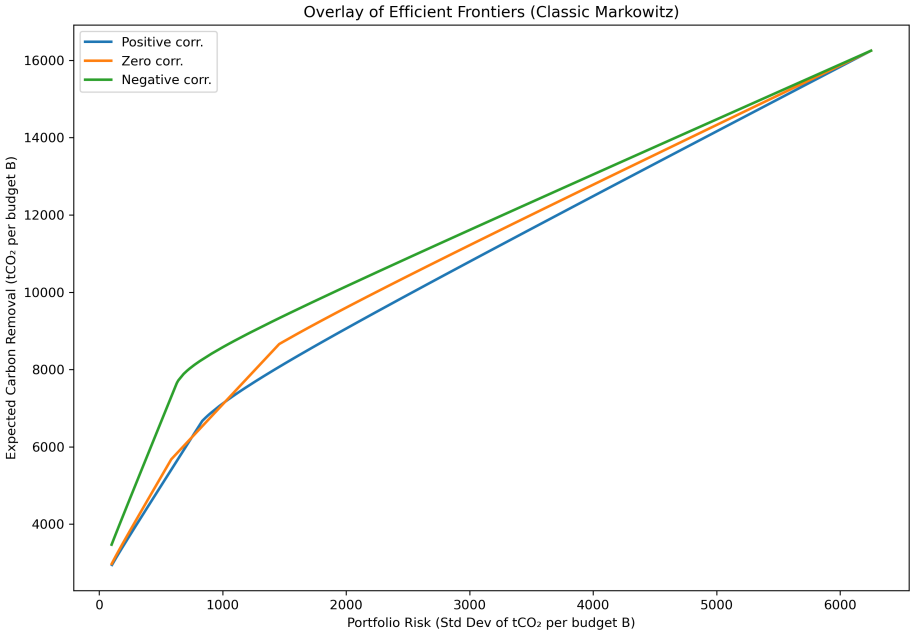

Figure S14. Efficient portfolios along the frontier in each correlation scenario. DACS dominates at low risk, biochar at intermediate risk, and forests at high expected return. Negative correlation induces stronger mixing of forest and DAC.

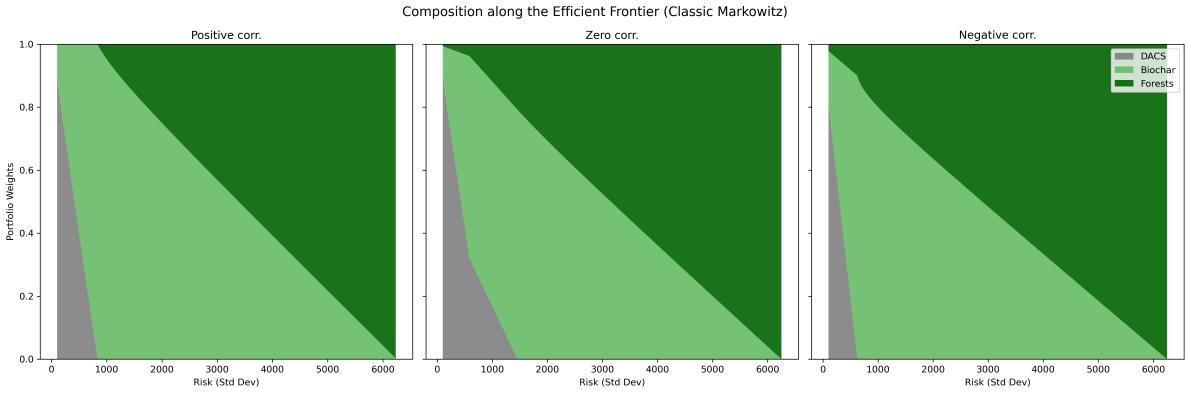

Figure S15. Composition of efficient portfolios along the frontier in each correlation scenario. DACS dominates at low risk, biochar at intermediate risk, and forests at high expected return. Negative correlation induces stronger mixing of forest and DAC. Diversification has greater benefits (higher expected return for given risk) when there is a negative correlation.

References

1. Bi, W. *et al.* A global 0.05° dataset for gross primary production of sunlit and shaded vegetation canopies from 1992 to 2020. *Sci. Data* **9**, 213, DOI: 10.1038/s41597-022-01309-2 (2022).
2. Shen, Y., Prentice, I. C. & Harrison, S. P. Investigation of factors that affect post-fire recovery of photosynthetic activity at global scale. *Ecol. Indic.* **171**, 113206, DOI: 10.1016/j.ecolind.2024.113206 (2025).
3. Alcalde, J. *et al.* Estimating geological CO2 storage security to deliver on climate mitigation. *Nat. Commun.* **9**, 2201, DOI: 10.1038/s41467-018-04423-1 (2018).
4. Weller, Z. D., Hamburg, S. P. & von Fischer, J. C. A national estimate of methane leakage from pipeline mains in natural gas local distribution systems. *Environ. Sci. & Technol.* **54**, 8958–8967, DOI: 10.1021/acs.est.0c00437 (2020).
5. Lamb, B. K. *et al.* Direct measurements show decreasing methane emissions from natural gas local distribution systems in the united states. *Environ. Sci. & Technol.* **49**, 5161–5169, DOI: 10.1021/es505116p (2015).
6. Markowitz, H. Portfolio selection. *The J. Finance* **7**, 77–91, DOI: 10.1111/j.1540-6261.1952.tb01525.x (1952).
7. Markowitz, H. *Portfolio Selection: Efficient Diversification of Investments* (John Wiley & Sons, 1959).