Supplementary Information

The Carbon at Risk measure to manage carbon removal risk and guide effective portfolio construction

Bali Lee[1,2], **Nick Gogerty**[3], **Tom Bearpark**[5,7], **Oleksandr Kit**[4], **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".

Contents

List of Figures

List of Tables

1 Value at Risk (VaR) and its adaptation to Carbon at Risk (CaR)

Carbon at Risk adapts the concept of Value at Risk (VaR) from financial risk management. VaR answers the question: *how much money could I lose?* For a confidence level α, VaR is the maximum loss not exceeded with probability α: $\Pr(L \leq \mathrm{VaR}_\alpha) = \alpha$, where L is the portfolio loss over a specified horizon. Figure S1 illustrates the concept for a hypothetical financial portfolio with normally distributed returns. The shaded areas show the lower tail at three confidence levels (90%, 95%, 98%), with higher confidence corresponding to larger potential losses.

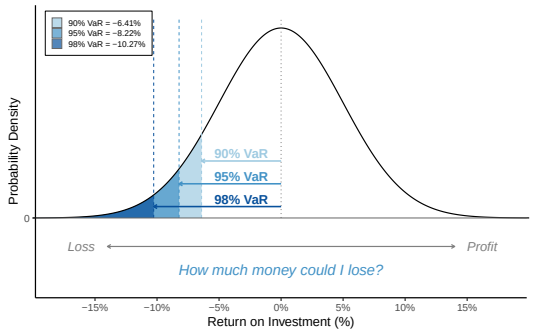

Figure S1. Value at Risk illustrated for a financial portfolio. The probability density of portfolio returns for an illustrative example portfolio is shown, with shaded regions indicating the lower tail at 90%, 95%, and 98% confidence levels. VaR at each level is the return threshold below which losses occur with the corresponding tail probability. For example, the 95% VaR of -8.22% means there is a 5% chance of losing more than 8.22% of the portfolio's value.

Just as VaR is used for different reasons for different stakeholders within financial markets, CaR is designed to serve corresponding roles within the carbon removal ecosystem. Table S1 maps ten stakeholder archetypes to their financial-market analogues, summarising how each would use CaR in practice.

CaR Stakeholder	VaR Analogue	How CaR is used
Project Developer	Trading Desk	Quantify expected carbon losses across time horizons; optimise project design; negotiate evidence-based buffer contributions
Credit Buyer	Portfolio Manager	Construct diversified removal portfolios with known risk budgets; match removal durability to emission profiles
Insurer / Reinsurer	Underwriter	Develop actuarially sound premium pricing based on measured CaR scores; create performance-based incentives
Financial Regulator	Basel Committee	Set dynamic, evidence-based buffer requirements and compliance thresholds calibrated to project-level CaR
Standards Body	Credit Rating Agency	Establish CaR-based credit tiers (e.g. AAA: CaR < 10 kg/t; A: CaR < 100 kg/t) for standardised comparison
Investor / Bank	PE Fund	Create investable carbon removal funds with defined risk–return profiles comparable to infrastructure debt
Credit Trader	Market Maker	Price credits on a risk-adjusted basis; establish replacement ratios across CaR tiers; develop derivative products
Verifier / Auditor	External Auditor	Verify declared CaR scores against observed performance data; provide independent quantitative assurance
Policymaker	Central Bank	Calibrate procurement incentives and tax credits to CaR scores; design NDC-aligned removal policy
Researcher / MRV	Quant Analyst	Refine CaR models, validate assumptions, discover cross-project risk correlations

Table S1. Stakeholder mapping: CaR roles and their VaR analogues in financial markets. Each stakeholder uses CaR in a manner that mirrors how the corresponding financial-market participant uses VaR.

Ref.[1] define Value at Risk as: "The maximum potential loss in value of a portfolio due to adverse market movements,

for a given probability." CaR is obtained from VaR by a direct substitution of the underlying quantities. Table S2 makes this correspondence explicit.

VaR concept		CaR concept
Loss in portfolio value	$\longrightarrow$	Shortfall in carbon removal
Financial portfolio	$\longrightarrow$	Removal project or portfolio
Adverse market movements	$\longrightarrow$	Adverse project performance

Table S2. Mapping between VaR and CaR definitions. CaR inherits the quantile-based structure of VaR but replaces financial loss with physical carbon shortfall.

2 Buffer Requirements and the Risk Fraction

This section provides the details underlying the buffer interpretation of CaR presented in the main text.

2.1 The buffer identity

A buyer who must deliver Q^* tonnes with α-level confidence contracts $Q = Q^* + B$ tonnes, where B is the buffer. The buffer is chosen so that the $(1-\alpha)$-quantile of actual delivery equals the target:

$$p_{1-\alpha}(Q) = Q^*.$$

Since $\mathrm{CaR}_\alpha = Q - p_{1-\alpha}(Q)$, we have:

$$B = Q - Q^* = Q - p_{1-\alpha}(Q) = \mathrm{CaR}_\alpha.$$

That is, the required buffer equals CaR when both are computed at the contracted scale Q.

A natural concern is that buffer tonnes will also fail, requiring "buffer on buffer." This concern is unfounded because the quantile $p_{1-\alpha}$ is computed on *all* Q contracted tonnes — target and buffer alike. The distribution F_D is the distribution of total delivery from the *entire* portfolio of Q tonnes. Nothing is excluded, so no additional correction is needed.

2.2 The risk fraction and the geometric-series interpretation

Define the **risk fraction** $\varphi_\alpha \equiv \mathrm{CaR}_\alpha / Q$. Then:

$$\begin{aligned} Q - \varphi_\alpha Q &= Q^*, \\ Q(1-\varphi_\alpha) &= Q^*, \\ Q &= \frac{Q^*}{1-\varphi_\alpha}. \end{aligned}$$

The factor $1/(1-\varphi_\alpha)$ has a geometric-series expansion:

$$\frac{1}{1-\varphi_\alpha} = 1 + \varphi_\alpha + \varphi_\alpha^2 + \varphi_\alpha^3 + \cdots$$

Each term accounts for the next round of buffer failure: the first term is the target, the second is the buffer for the target, the third is the buffer for the buffer, and so on. The CaR-based buffer formula automatically sums all these rounds.

2.3 Numerical example

Here, we provide an illustrative numerical example: a buyer needs to deliver $Q^* = 1{,}000$ tonnes with 95% confidence. Two technologies are considered: a low-risk technology with $\varphi_{95} = 0.03$ and a high-risk technology with $\varphi_{95} = 0.45$.

	DACCS ($\varphi_{95} = 0.03$)	**CA Forest** ($\varphi_{95} = 0.45$)
Target Q^* (tonnes)	1,000	1,000
Naive buffer ($\varphi \times Q^*$)	30	450
Correct $Q = Q^*/(1-\varphi)$	1,031	1,818
Correct buffer	31	818
Naive undercount	1 ($\sim$3%)	368 ($\sim$45%)

For low-risk technologies the naive approach (computing CaR at the target scale Q^* and using it directly as the buffer) is nearly exact. For high-risk technologies it underestimates the required buffer by up to 45%, because it fails to account for the fact that buffer tonnes also experience the φ failure rate.

2.4 Does φ change when the portfolio scales?

The formula $Q = Q^*/(1-\varphi_\alpha)$ is exact for any given value of φ_α. The subtlety is that adding buffer projects enlarges the portfolio, and φ_α may itself change with portfolio size. In general, φ_α depends on the number of projects K and the correlation structure ρ across projects. To see this, consider K identical projects each delivering q tonnes with survival probability p, so $Q = Kq$. Under a normal approximation (valid for large K):

$$\mathrm{CaR}_{95} \approx Kq(1-p) + 1.645 \cdot \mathrm{SD}(D),$$

where $\mathrm{SD}(D)$ depends on the correlation structure.

Independent projects ($\rho = 0$).

$$\varphi(K) = (1-p) + \frac{1.645\sqrt{p(1-p)}}{\sqrt{K}}.$$

The risk fraction decreases in K — this is diversification. As $K \to \infty$, $\varphi \to (1-p)$: only the expected loss remains.

Correlated projects ($\rho > 0$).

$$\varphi(K) = (1-p) + 1.645\sqrt{p(1-p)} \cdot \sqrt{\frac{1-\rho}{K} + \rho}.$$

The risk fraction still decreases in K but is bounded below by $\varphi(\infty) = (1-p) + 1.645\sqrt{p(1-p)\rho}$. The idiosyncratic component $(1-\rho)/K$ vanishes with diversification; the systematic component ρ is irreducible.

Perfect correlation ($\rho = 1$). φ is exactly constant for all K. The closed-form buffer formula is exact.

In practice, φ is approximately constant when: (i) ρ is high; (ii) K is large; or (iii) p is high (low-risk technology, small buffer, small change in K). The approximation is roughest for independent, high-risk projects at small portfolio sizes. In such cases, the formula overestimates the required buffer (a conservative error), and the exact solution should be obtained numerically by solving $p_{1-\alpha}(Q) = Q^*$.

2.5 Effective cost

In general, the effective cost per tonne delivered at confidence level α is defined as the total portfolio cost divided by the α-guaranteed delivery:

$$c_{\text{eff},\alpha} = \frac{C}{p_{1-\alpha}},$$

where C is total portfolio cost and $p_{1-\alpha}$ is the $(1-\alpha)$-quantile of realised delivery. This definition applies to any portfolio, including multi-technology mixes where there is no single sticker price.

For the single-technology case, $C = cQ$ where c is the sticker price per tonne contracted. Since the buyer must contract $Q = Q^*/(1-\varphi_\alpha)$ tonnes to deliver Q^* at confidence α, and $p_{1-\alpha} = Q - \text{CaR}_\alpha = Q(1-\varphi_\alpha) = Q^*$, the general formula reduces to:

$$c_{\text{eff},\alpha} = \frac{cQ}{Q(1-\varphi_\alpha)} = \frac{c}{1-\varphi_\alpha}.$$

The denominator $(1-\varphi_\alpha)$ accounts for all rounds of buffer failure. As with the buffer formula, $c_{\text{eff},\alpha}$ is conditional on a given portfolio composition, since φ_α depends on portfolio size K and correlation ρ.

3 Forest Carbon Fire Risk

This section presents the data inputs, parameters, and supplementary analysis supporting the forest carbon CaR estimates in the main text.

3.1 Data and Parameters

Table S3 summarises the default simulation parameters.

Parameter	Value
Annual climate change-induced burn rate increase	0.5%
Number of Monte Carlo simulations	1,000
Annual regrowth rate, California	2.0%/yr
Annual regrowth rate, Mato Grosso	3.0%/yr
Annual regrowth rate, Papua	1.5%/yr

Table S3. Parameters for the forest CaR simulation

The data sources and simulation procedure are described in Methods.
Regional fire history for the three focal regions are shown in Fig. S2.

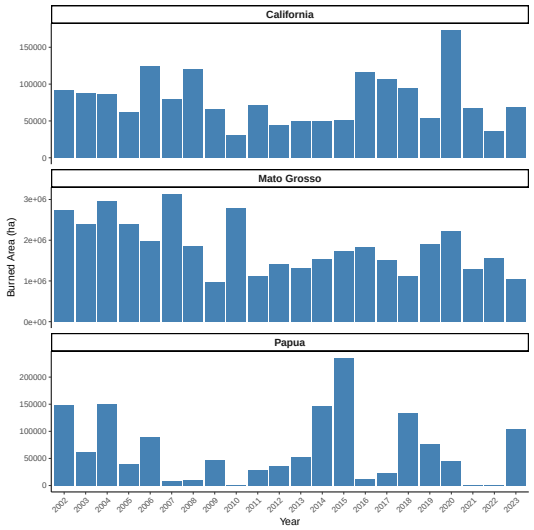

Figure S2. Annual burned area (hectares) for three focal regions, 2002–2023. Data from EFFIS.

Caveats. Fires reported by EFFIS 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, which may overestimate fire hazard. This problem can be mitigated by applying spatio-temporal clustering algorithms, which are not part of this study. For simplification, this analysis does not take into account project management risks or the political and socio-economic landscape. We provide a starting point for more detailed analysis in the future. Further, climate change is likely to cause heterogenous changes in fire risk, a feature not captured in our analysis.

3.2 CaR Curve Dynamics: Analytic Characterisation

The CaR curves in main-text Fig. 2b exhibit three distinct phases whose origins can be understood through a deterministic state-space approximation. Let $U_t \in [0,1]$ denote the fraction of carbon remaining intact at time t. Each period, the state evolves via two competing processes: burning of intact forest at rate $\tilde{\beta}_t$ and regrowth of burned forest at rate r. These operate simultaneously on the previous period's state, giving:

$$U_t = U_{t-1}(1-\tilde{\beta}_t) + r(1-U_{t-1}), \tag{1}$$

where $\tilde{\beta}_t = \beta \cdot (1+\gamma(t-1))$ is the climate-adjusted burn rate, β is drawn from the empirical distribution, r is the regrowth rate, and γ is the annual climate trend.

Stationary equilibrium ($\gamma = 0$). When the burn rate is constant at its mean $\bar{\beta}$, setting $U_t = U_{t-1} = U^*$ yields a fixed point:

$$U^* = \frac{r}{r+\bar{\beta}}, \qquad \mathbb{E}[L^*] = 1 - U^* = \frac{\bar{\beta}}{r+\bar{\beta}}, \tag{2}$$

where $L^* = 1 - U^*$ is the expected equilibrium carbon loss. For California ($r = 2.0\%$/yr, $\bar{\beta} = 3.6\%$/yr), $\mathbb{E}[L^*] \approx 0.64$, corresponding to an expected loss of 643 kg/tonne.

Moving equilibrium ($\gamma > 0$). Under a positive climate trend, the equilibrium becomes time-varying:

$$\mathbb{E}[L_t^*] = \frac{\bar{\beta}(1+\gamma t)}{r+\bar{\beta}(1+\gamma t)}. \tag{3}$$

As $t \to \infty$, $\mathbb{E}[L_t^*] \to 1$. The system perpetually "chases" a shifting target, producing the slow upward drift visible in the CaR curves at long horizons.

Three phases. Together, these results explain the characteristic shape of the CaR curves:

1. **Transient** (∼0–30 yr): Starting from $U_0 = 1$, burning greatly exceeds regrowth, so carbon loss accumulates rapidly.
2. **Quasi-equilibrium** (∼30–80 yr): The system approaches U^*; burning approximately balances regrowth, and the CaR curve flattens.
3. **Climate drift** (∼80+ yr): The increasing burn rate shifts the equilibrium upward, causing CaR to resume a slow rise.

The 95% CaR lies above the expected loss by an amount that reflects the dispersion of the burn-rate distribution, approximated as $\text{CaR}_{95} \approx \mathbb{E}[L_T] + 1.645 \cdot \sigma_T$, where σ_T is the standard deviation of cumulative loss at horizon T. Note that this is a Gaussian approximation that may not hold in the tails, particularly at short horizons where the loss distribution is skewed; the Monte Carlo CaR estimates in the main text do not rely on this approximation.

3.3 Sensitivity to the Fire Distribution Assumption

The baseline CaR estimates in the main text are obtained by resampling annual burn fractions from the 22-year EFFIS record (2002–2023) for each region. Because CaR is a tail quantity—the 95th percentile of cumulative carbon loss—it is natural to ask whether restricting draws to the 22 observed values understates tail risk.

Peaks-over-threshold approach. We construct an alternative fire-risk model using the Pickands–Balkema–de Haan theorem, which states that exceedances above a sufficiently high threshold u follow a Generalised Pareto Distribution (GPD):

$$G_{\xi,\sigma}(x) = 1 - \left(1 + \xi\, x/\sigma\right)^{-1/\xi}, \qquad x > 0,\ \ \sigma > 0. \tag{4}$$

We set u equal to the median of the California burn fractions, yielding $n_u = 11$ exceedances. Maximum likelihood estimation gives $\hat{\sigma} = 0.012$ and $\hat{\xi} = 0.67$, indicating a heavy-tailed (Pareto-type) distribution. A spliced model draws from the empirical distribution below u and from the fitted GPD above u, preserving the bulk while allowing the right tail to extend beyond the historic maximum (Fig. S3).

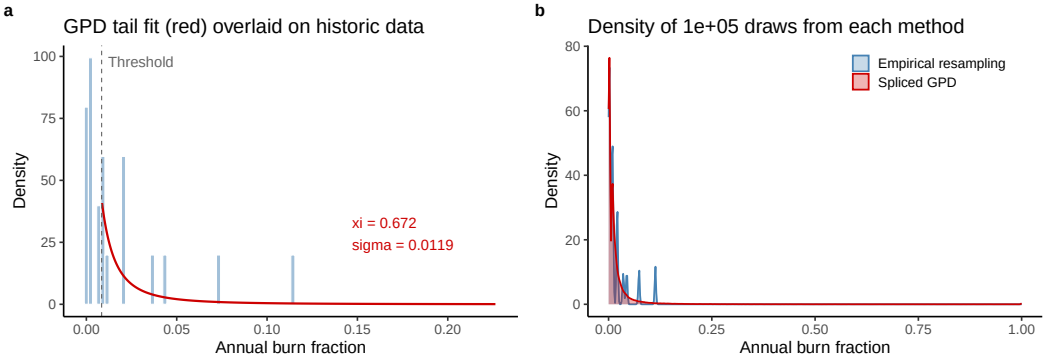

Figure S3. Spliced GPD tail model for California burn fractions. a, Histogram of 22 annual burn fractions (2002–2023) with the fitted GPD density (red) above the median threshold (dashed). **b**, Density of 100,000 draws from the empirical bootstrap (blue) and the spliced GPD model (red), showing agreement in the bulk and divergence in the right tail.

Impact on CaR. We feed both sampling methods into the same Monte Carlo simulation engine ($N = 5{,}000$ paths, $\gamma = 0.5\%$/yr, regrowth $r = 2.0\%$/yr). Figure S4 compares the resulting 95% CaR curves. The GPD-based CaR is approximately 20% higher than the empirical estimate at a 100-year horizon, with larger gaps at intermediate horizons ($T \approx 10$–30 yr), where a single extreme fire year can dominate the 5th-percentile path. At long horizons, regrowth partially offsets extreme burns, narrowing the gap. The baseline empirical CaR estimates reported in the main text are therefore conservative with respect to heavy-tailed fire risk.

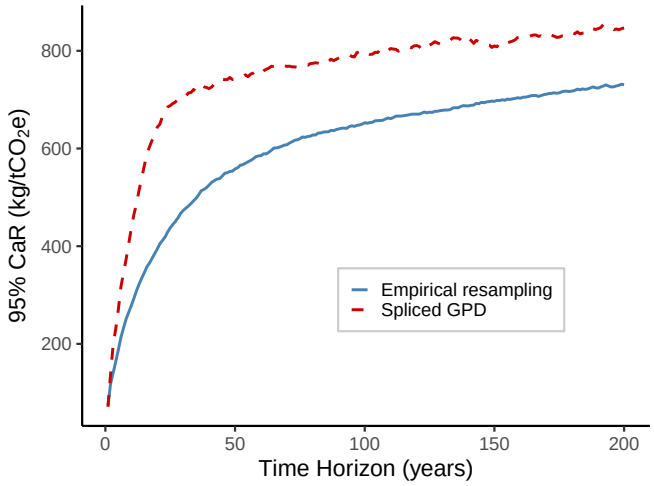

Figure S4. Sensitivity of 95% CaR to the fire distribution assumption (California, single project). The empirical resampling baseline (solid blue) is compared with the spliced GPD model (dashed red). The GPD model, which allows burn fractions to exceed the historic maximum, produces CaR estimates 20–60% higher.

Caveats. With only $n_u = 11$ exceedances, the GPD parameter estimates carry substantial uncertainty; the point estimate $\hat{\xi} = 0.67$ is suggestive of heavy tails but not definitive. The threshold choice ($u =$ median) is pragmatic rather than optimised. Furthermore, the spliced model assumes a time-invariant tail shape, whereas climate change may alter ξ over time. These results should therefore be interpreted as a directional sensitivity check rather than a definitive alternative calibration.

3.4 Sensitivity to the Climate Trend Parameter (γ)

The baseline CaR estimates in the main text assume a climate-driven annual increase in burn rate of $\gamma = 0.5\%$/yr. This value is broadly consistent with recent estimates of climate-attributable trends in global burned area[2], though regional trends vary substantially. Fig. S5 reports CaR curves for $\gamma \in \{0\%, 0.5\%, 1.0\%\}$ across all three study regions. At a 200-year horizon, the choice of γ has a large effect on CaR: for California, CaR_{95} ranges from 56% ($\gamma = 0$) to 82% ($\gamma = 1\%$/yr). The sensitivity is smaller in Papua, where the low baseline fire risk limits the climate amplification effect. These results underscore the importance of the climate trend assumption and motivate reporting CaR under a range of scenarios.

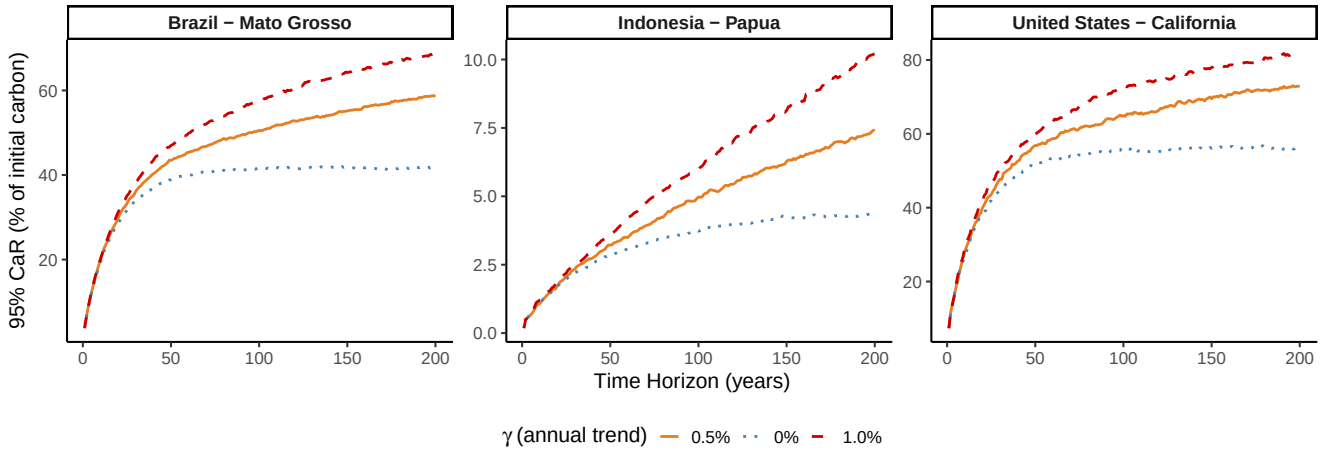

Figure S5. Sensitivity of 95% CaR to the climate trend parameter γ. CaR curves for three values of γ (0%, 0.5%, 1.0% annual increase in burn rate) across three regions ($N = 1{,}000$ Monte Carlo simulations).

3.5 Sensitivity to the Regrowth Rate

The baseline regrowth rate for California is $r = 2.0\%$/yr, calibrated from post-fire recovery data[3] with a literature-based adjustment (see Methods). Fig. S6 reports CaR curves when the regrowth rate is scaled by factors of 0.5, 1.0, and 1.5. At a 200-year horizon, halving the regrowth rate increases CaR_{95} from 73% to 84%, while increasing it by 50% reduces CaR to 65%. The sensitivity is substantial but smaller than the effect of γ, reflecting the fact that regrowth operates as a steady counterweight to fire losses while the climate trend compounds over time.

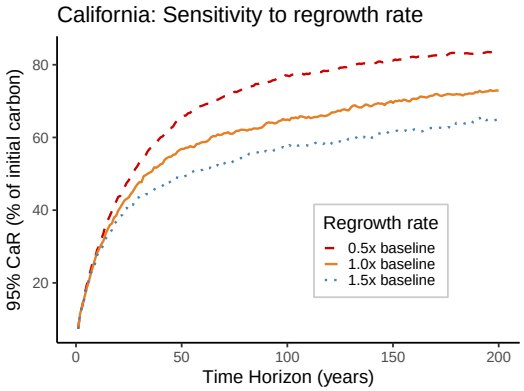

Figure S6. Sensitivity of 95% CaR to the regrowth rate (California). CaR curves for three regrowth rate scalings (0.5×, 1.0×, 1.5× baseline rate of 2.0%/yr; $N = 1{,}000$ Monte Carlo simulations).

3.6 Diversification and Portfolio Convergence under Correlation

Main-text Fig. 2d,e summarise the diversification benefit of holding a portfolio of K projects within a single region, as a function of inter-project correlation ρ. This section provides additional detail on the mechanics and convergence properties.

Impact of correlation on the CaR gap. Figure S7 displays the full 200-year CaR trajectories for a single project ($K = 1$) and a 100-project portfolio ($K = 100$) at three correlation levels. Under independence ($\rho = 0$), the portfolio CaR (blue dashed) converges closely to the expected loss (black solid), reflecting classical diversification. At moderate correlation ($\rho = 0.5$), the portfolio CaR remains substantially above the expected loss: the common-factor component of fire risk cannot be diversified away. At perfect correlation ($\rho = 1$), the single-project and portfolio CaR curves coincide, confirming that diversification provides no benefit when all projects burn together. These trajectories illustrate how the irreducible variance floor $\rho\sigma^2$ manifests in the time-domain: it widens the gap between the portfolio CaR and the expected loss at every horizon.

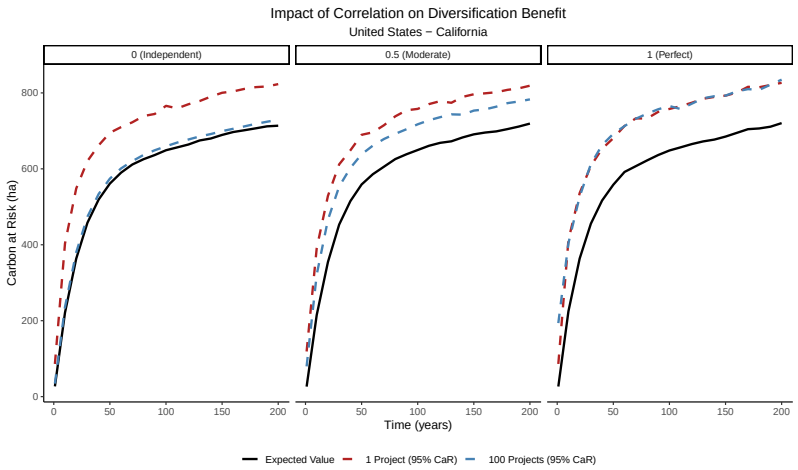

Figure S7. Impact of inter-project correlation on diversification benefit (California). Each panel shows the expected cumulative carbon loss (black solid), 95% CaR for a single project (red dashed), and 95% CaR for a 100-project portfolio (blue dashed) over 200 years, at correlation $\rho = 0$ (independent), $\rho = 0.5$ (moderate), and $\rho = 1$ (perfect). Under independence the portfolio CaR closely tracks the expected loss; as ρ increases, the portfolio CaR converges toward the single-project CaR.

Convergence in K across correlation levels. Figure S8 traces the 100-year 95% CaR as a function of portfolio size K for five correlation levels ($\rho \in \{0, 0.05, 0.1, 0.25, 0.5\}$). Under independence ($\rho = 0$), the CaR declines steadily toward the expected loss (dashed line) as K increases. For any $\rho > 0$, the CaR curve flattens well above the expected loss, with the plateau rising in ρ. Diminishing returns set in rapidly: most of the achievable diversification benefit is realised by $K \approx 20$–50 projects, regardless of correlation. This has practical implications for portfolio design: beyond a moderate number of within-region projects, further risk reduction requires spreading across regions with lower inter-regional correlation.

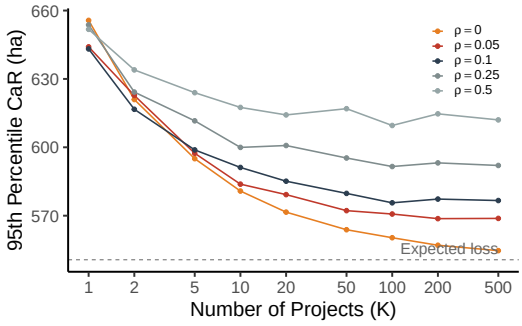

Figure S8. Portfolio CaR convergence as a function of the number of projects K and inter-project correlation ρ (California, 100-year horizon). The 95% CaR is shown for $K \in \{1, 2, 5, 10, 20, 50, 100, 200, 500\}$ at five correlation levels. The dashed line marks the expected loss. Under independence ($\rho = 0$), CaR converges to the expected loss; for $\rho > 0$, CaR plateaus above the expected loss at a floor that increases with ρ. $N = 1{,}000$ Monte Carlo simulations per (K, ρ) pair.

4 Geological Storage: Extended Horizons

The main text presents CaR estimates for geological CO_2 storage at a 200-year horizon. This section extends the analysis to 1,000 and 10,000-year horizons using the Storage Security Calculator of[4], with 10,000 Monte Carlo realisations per scenario. At longer horizons, absolute leakage increases but trapping mechanisms (mineral, solubility, residual) progressively immobilise a larger fraction of injected CO_2.

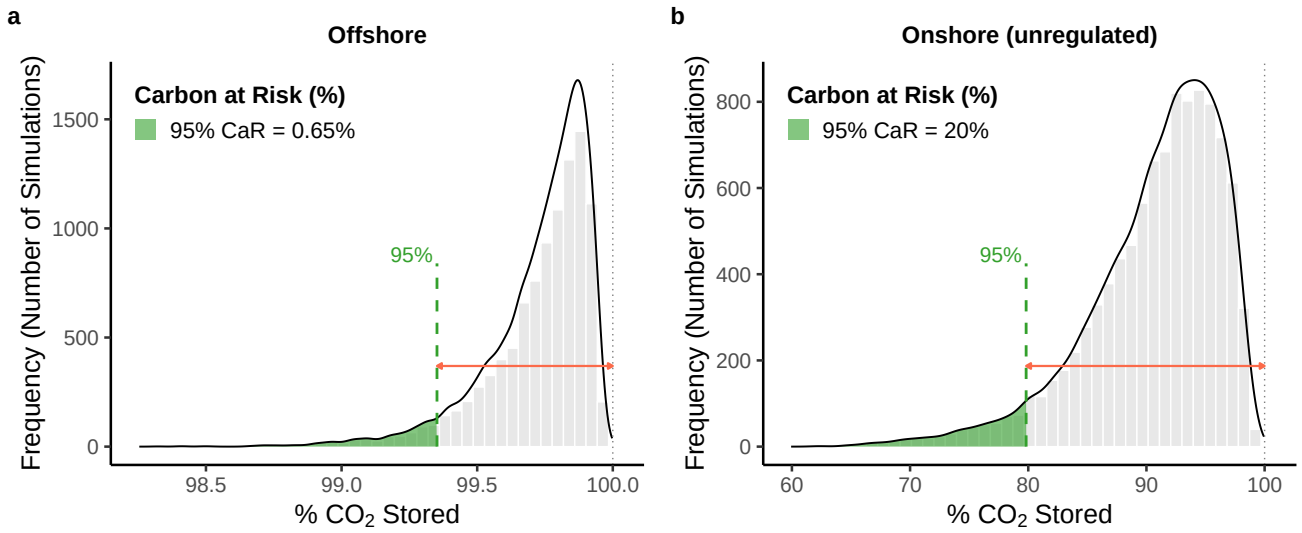

Figure S9. DACCS Carbon at Risk at a 1,000-year horizon. Monte Carlo simulation (10,000 realisations) of CO_2 storage retention using the Storage Security Calculator of[4]. (a) Well-regulated offshore scenario: 95% CaR = 0.65%. (b) Poorly regulated onshore scenario: 95% CaR = 20.2%. The two-orders-of-magnitude gap persists at longer horizons, though absolute losses increase as trapping mechanisms are not yet fully effective.

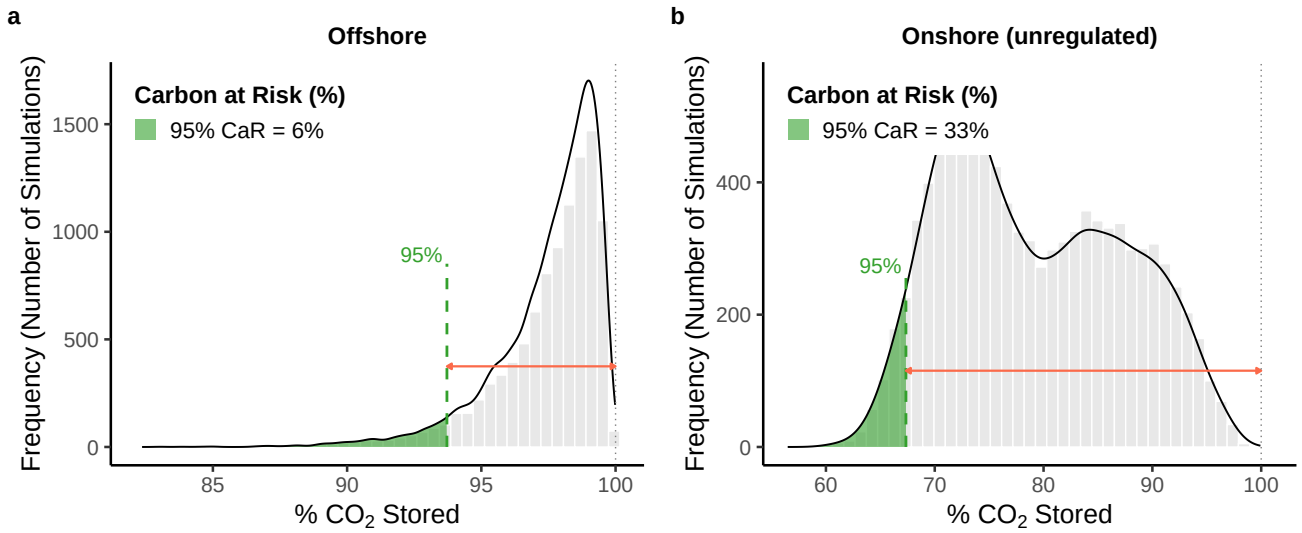

Figure S10. DACCS Carbon at Risk at a 10,000-year horizon. Monte Carlo simulation (10,000 realisations) as in the previous figure but over 10,000 years. (a) Well-regulated offshore scenario: 95% CaR = 6.3%. (b) Poorly regulated onshore scenario: 95% CaR = 32.6%. At this horizon, even the well-regulated offshore scenario shows non-trivial leakage, underscoring the importance of monitoring and trapping efficiency assumptions over geological timescales.

5 Portfolio Optimisation under Correlation

This section derives the cost-minimising portfolio of heterogeneous carbon-removal projects when within-technology correlation is non-zero, extending the classical mean–variance framework[5] to the Bernoulli delivery setting of the main text.

5.1 Setup and Variance Structure

Consider a portfolio of n_D DACCS projects and n_F forest projects, each contracted to deliver q tonnes of CO_2 if successful. Project j of technology $i \in \{D,F\}$ is Bernoulli: it delivers q with probability p_i and zero otherwise. Within-technology pairwise correlation is ρ_i; the between-technology correlation is ρ_{between}.

Portfolio mean.

$$\mu = q\big(n_D p_D + n_F p_F\big). \tag{5}$$

Portfolio variance.

$$\sigma^2 = \sigma_D^2 + \sigma_F^2 + 2\,\mathrm{Cov}_{DF}, \tag{6}$$

where the within-technology variance is

$$\sigma_i^2 = q^2\, p_i(1-p_i)\big[n_i + n_i(n_i-1)\rho_i\big], \tag{7}$$

and the between-technology covariance is

$$\mathrm{Cov}_{DF} = n_D\, n_F\, \rho_{\text{between}}\, q^2 \sqrt{p_D(1-p_D)\, p_F(1-p_F)}. \tag{8}$$

Under a normal approximation, the 5th percentile of portfolio delivery is

$$p_5 = \mu + z\sigma, \qquad z = \Phi^{-1}(0.05) \approx -1.645, \tag{9}$$

and the Carbon at Risk at 95% confidence is

$$\mathrm{CaR}_{95} = (n_D + n_F)\, q - p_5. \tag{10}$$

5.2 Lagrangian Formulation of the Three Procurement Rules

We present the formal optimisation problem for each of the three procurement rules introduced in Section 4. In each case the buyer chooses $(n_D, n_F) \geq 0$ to minimise total cost $C = c_D n_D + c_F n_F$.

Rule 1 (cheapest tonne). The buyer targets contracted volume:

$$\min_{n_D, n_F \geq 0} C \quad \text{s.t.} \quad Q = q(n_D + n_F) \geq T. \tag{11}$$

The Lagrangian is $\mathscr{L} = C - \lambda[q(n_D + n_F) - T]$. The first-order conditions require $c_i = \lambda q$ for each active technology. Since $c_F < c_D$, only the forest FOC can bind: the solution is always a corner at all forest, $n_F^* = \lceil T/q \rceil$, $n_D^* = 0$.

Rule 2 (only permanent). The buyer restricts to DACCS ($n_F = 0$) and selects the minimum number of DACCS projects such that $p_5 \geq T$:

$$\min_{n_D \geq 0} c_D n_D \quad \text{s.t.} \quad n_F = 0,\ \ p_5(n_D, 0) \geq T. \tag{12}$$

This is a one-dimensional problem. Because p_5 is increasing in n_D, the solution is $n_D^* = \min\{n_D : p_5(n_D, 0) \geq T\}$. No optimisation over the technology mix occurs; the rule simply sizes the DACCS portfolio to meet the delivery guarantee.

Rule 3 (CaR-constrained). The buyer minimises cost subject to a quantile delivery target:

$$\min_{n_D, n_F \geq 0} C \quad \text{s.t.} \quad p_5(n_D, n_F) \geq T. \tag{13}$$

Relaxing to continuous n_i and writing $p_5 = \mu + z\sigma$ with $z = \Phi^{-1}(0.05) \approx -1.645$, the Lagrangian is

$$\mathscr{L} = c_D n_D + c_F n_F + \lambda\big(T - \mu - z\sigma\big), \qquad \lambda \geq 0. \tag{14}$$

The first-order conditions yield the equal bang-per-buck condition

$$\frac{MV_D}{c_D} = \frac{MV_F}{c_F}, \tag{15}$$

where the marginal value of an additional project of technology i decomposes as

$$MV_i = q\,p_i - \frac{z}{2\sigma}\frac{\partial \sigma^2}{\partial n_i}. \tag{16}$$

The first term is the expected carbon contribution; the second is the risk penalty, the shadow cost of the additional variance that the project introduces into the portfolio.

5.3 Corner vs Interior Solutions

In the Bernoulli model the marginal variance contribution is

$$\frac{\partial \sigma^2}{\partial n_i} = q^2\,p_i(1-p_i)\big[1+(2n_i-1)\rho_i\big]. \tag{17}$$

Independence ($\rho_i = 0$). The risk penalty in (16) is constant in n_i, so the technology with the higher marginal-value-to-cost ratio dominates everywhere. The solution is a corner: the buyer procures from a single technology. Rules 1 and 3 therefore give the same answer (all forest) under independence.

Correlation ($\rho_i > 0$). The risk penalty is increasing in n_i, reflecting diminishing returns to concentration in one technology. At some portfolio size the marginal value-to-cost ratio of the initially dominant technology falls below that of the other, and an interior solution emerges: mixing is optimal. An interior solution is more likely when (i) within-technology correlation ρ_i is high, or (ii) the cost gap c_D/c_F is moderate. This is the formal basis for why only Rule 3 diversifies under correlation, while Rules 1 and 2 always produce corner solutions.

5.4 Effective Price

The effective price per tonne delivered with α-level confidence is

$$c_{\text{eff},\alpha} = \frac{c}{1-\varphi_\alpha}, \qquad \varphi_\alpha = \frac{\text{CaR}_\alpha}{Q}, \tag{18}$$

where c is the sticker price per tonne and φ_α is the risk fraction—the share of contracted volume at risk at confidence level α.

5.5 Calibration

Table S4 summarises the illustrative calibration used in the main text.

Table S4. Portfolio optimisation calibration parameters.

Parameter	DACCS	Forest
Survival probability p_i	0.99	0.4
Cost per project c_i	\$500	\$40
Within-technology correlation ρ_i	0.5	0.2

Between-technology correlation $\rho_{\text{between}} = 0$. Each project delivers $q = 500$ tCO_2 if successful. Target: $p_5 \geq 30{,}000$ tCO_2. The DACCS within-technology correlation reflects shared engineering and supply-chain risk (illustrative). The forest within-technology correlation is calibrated from MODIS burned-area data (see Section 3).

5.6 Additional Portfolio Figures

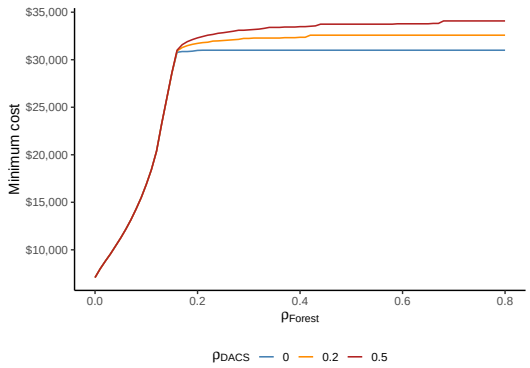

Figure S11. Minimum portfolio cost as a function of within-forest correlation. For three values of ρ_{DACCS}, the figure shows the minimum total cost to meet the $p_5 = 30{,}000$ tCO_2 target as ρ_{Forest} increases. Higher forest correlation raises portfolio cost because more projects are needed to offset the reduced diversification benefit.

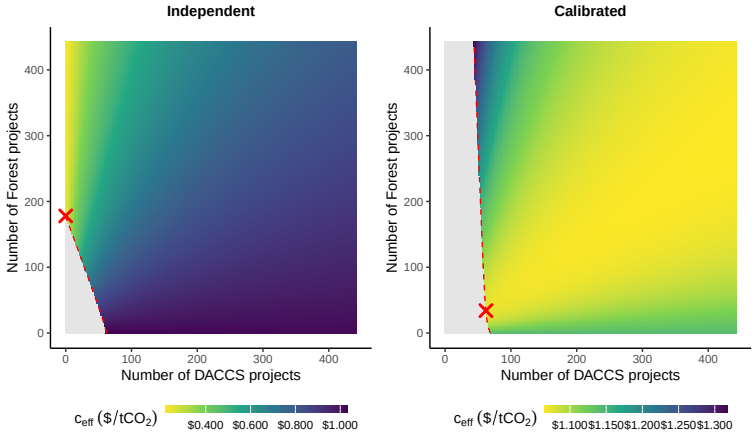

Figure S12. Effective price per tonne delivered. Heatmap over $(n_{\text{DACCS}}, n_{\text{Forest}})$ coloured by the effective price $c_{\text{eff}} = \text{cost}/p_5$, under independence (left) and calibrated correlations (right). Grey region is infeasible ($p_5 < 30{,}000$ tCO_2). Red dashed contour marks the feasibility boundary; red cross marks the minimum-cost portfolio. Under both scenarios the optimum is an interior mix, but correlation shifts it toward a more DACCS-heavy composition.

5.7 Expected-Value Procurement Rule

In the main text, Rule 1 targets the cheapest contracted tonne ($Q \geq T$), which buys 60 forest projects for $2,400. A more sophisticated variant targets expected delivery ($\mu \geq T$), buying 150 forest projects for $6,000. Both rules yield the same qualitative outcome: all forest, corner solution. The expected-value variant over-purchases to compensate for expected failures, but still ignores tail risk. Under calibrated correlations ($\rho_F = 0.2$), the expected-value rule achieves $\mu = 30{,}000$ tCO_2 but $p_5 = 2{,}614$ tCO_2, far below the target. Like the $Q \geq T$ rule, it never diversifies across technologies because its constraint is linear in (n_D, n_F). The key results of Section 4 are unchanged: only the CaR-constrained rule (Rule 3) produces an interior mix under correlation.

References

1. Engle, R. F. & Manganelli, S. Value at risk models in finance. Working Paper Series 75, European Central Bank (2001).
2. Boer, M. M., Bowman, D. M. J. S., Bradstock, R. A. *et al.* Global burned area increasingly explained by climate change. *Nat. Clim. Chang.* **14**, 1186–1192, DOI: 10.1038/s41558-024-02140-w (2024).
3. Zang, Y., Quirk, B. & Shendryk, Y. A global dataset of forest regrowth following wildfires. *Sci. Data* **11**, 1052, DOI: 10.1038/s41597-024-03896-8 (2024).
4. 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).
5. Markowitz, H. Portfolio selection. *The J. Finance* **7**, 77–91, DOI: 10.1111/j.1540-6261.1952.tb01525.x (1952).