SUPPLEMENTARY MATERIAL for "Uncertainty quantification for computer models with spatial output using calibration-optimal bases"

S1 Idealised example

The spatial idealised example $f(\boldsymbol{\theta})$, introduced in Section 3.1, gives output over a 10×10 field, and has 6 input parameters defined on $[-1, 1]$:

$$f(\boldsymbol{\theta}) = 3(10\theta_2^2\boldsymbol{\varphi}_2 + 5\theta_3^2\boldsymbol{\varphi}_2 + (\theta_3 + 1.5\theta_1\theta_2)\boldsymbol{\varphi}_3 + 2\theta_2\boldsymbol{\varphi}_4 + \theta_3\theta_1\boldsymbol{\varphi}_5 + \theta_2\theta_1\boldsymbol{\varphi}_6 + \theta_2^3\boldsymbol{\varphi}_7 + (\theta_2 + \theta_3)^2\boldsymbol{\varphi}_8 + 2) + 1.5\pi_N(\theta_4, 0.2, 0.1^2)\boldsymbol{\varphi}_1\frac{\theta_5}{1.3 + \theta_6} + \Psi_{10\times10}(0, 0.05^2) \quad \text{(S1)}$$

for $\pi_N(\theta_4, 0.2, 0.1^2)$ the density function of the Normal distribution with mean 0.2 and variance 0.1^2, and where $\Psi_{10\times10}(0, 0.05^2)$ gives a sample from a Normal distribution with mean zero and variance 0.05^2, at each location in the 10×10 grid. Figure S1 shows $(\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_8)$, with $\boldsymbol{\varphi}_1$ giving the pattern most consistent with the observations, and $\boldsymbol{\varphi}_2$ the basis vector that dominates the ensemble. After evaluating $f(\boldsymbol{\theta})$, the output is vectorised so that $\ell = 100$.

We define $\boldsymbol{\theta}^*$ as

$$\boldsymbol{\theta}^* = (0.7, 0.01, 0.01, 0.25, 0.8, -0.9)$$

with the observed field, $\mathbf{z}$, given by adding a sample from $\mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{e}})$ to $f(\boldsymbol{\theta}^*)$, to represent observation error. We define $\boldsymbol{\Sigma}_{\mathbf{e}}$ using the squared exponential correlation function over the 10×10 grid, with the spatial coordinates denoted by $\mathbf{s}_i = (s_{i1}, s_{i2})$ for $i = 1, \dots, 100$. The i, j^{th} entry of 100×100 matrix $\boldsymbol{\Sigma}_{\mathbf{e}}$ is therefore

$$\boldsymbol{\Sigma}_{\mathbf{e}}^{ij} = \exp\{-(s_{i1} - s_{j1})^2 - (s_{i2} - s_{j2})^2\}. \quad \text{(S2)}$$

We model the discrepancy as

$$\boldsymbol{\eta} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\eta}}),$$

with the $(i, j)^{th}$ value of $\boldsymbol{\Sigma}_{\boldsymbol{\eta}}$ given by

$$\boldsymbol{\Sigma}_{\boldsymbol{\eta}}^{ij} = v_i v_j C(\mathbf{s}_i, \mathbf{s}_j) \quad \text{(S3)}$$

for variances v_i, v_j, and a correlation function $C(\cdot, \cdot)$ between locations $\mathbf{s}_i$ and $\mathbf{s}_j$. For $C(\cdot, \cdot)$, we again use the squared exponential correlation function, with the same correlation lengths as for $\boldsymbol{\Sigma}_{\mathbf{e}}$. We define v_i via

$$v_i = \begin{cases} 0.1 & \text{if } i \in S \\ 1 & \text{otherwise} \end{cases}$$

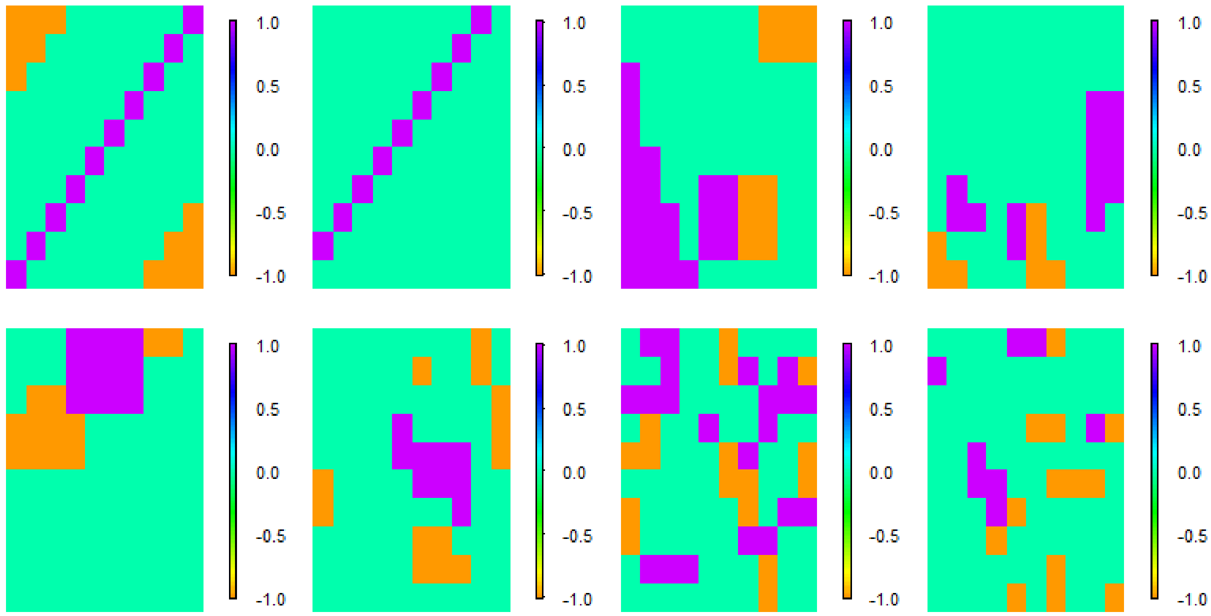

Figure S1: The 8 orthogonal basis vectors used in the definition of the idealised function, with $\boldsymbol{\varphi}_1$ shown in the top left, $\boldsymbol{\varphi}_2$ to the right of this, etc.

where the set S contains the grid boxes on the main diagonal, as we are more interested in finding fields with output consistent with the observations in this region of the output. Figure S2 shows the true NROY space given $\boldsymbol{\Sigma_e}$ and $\boldsymbol{\Sigma_\eta}$.

S1.1 Probabilistic calibration for the SVD basis

We performed probabilistic calibration following the Kennedy O'Hagan method, with fixed discrepancy as outlined in the previous section, uniform priors on the calibration parameters, and emulators as described in the main text.

The dotted lines on Figure S3 show the resulting densities when the truncated SVD basis, $\boldsymbol{\Gamma}_4$, is used to calibrate probabilistically on the field, for the input parameters $\theta_1, \ldots, \theta_5$, and the ratio $r = \theta_5/(1.3 + \theta_6)$. There are peaks of density away from the true values ($\boldsymbol{\theta}^*$, as shown by the red vertical lines), particularly for θ_3 and r. For θ_4, the parameter that controls the strength of the main diagonal, the posterior density is relatively flat across the entire range of θ_4.

We sample from these posteriors and run the idealised function at these samples, to evaluate whether calibration with the truncated SVD basis has highlighted a region of parameter space that is ‘close’ to $\mathbf{z}$. 16 samples are shown in Figure S4, demonstrating that the results suggest it is not possible to remove the off-diagonal pattern that was dominant in the ensemble.

S1.2 Calibration with the rotated bases

The solid lines in Figure S3 show the posterior distributions for $\theta_1, \ldots, \theta_5$ and r when the wave 1 rotated basis is used for probabilistic calibration, showing improvements (compared to the SVD basis) for θ_4 and r.

At wave 2, there are peaks of density at or near to the true parameter values for all but θ_5 and r, as shown by the solid lines in Figure S6. The dotted lines in this plot show the wave 3 posteriors. Although the peaks for θ_2, θ_3 and θ_4 are not as large, this wave offers an improvement for r (important for the strength of the main diagonal) and θ_1 (as this parameter has no effect on the main or off-diagonal, a flat posterior is more accurate). The averaged posterior samples in Figure 5 show that the posterior distributions at each wave have identified better regions of parameter space than previously, with the wave 3 samples being consistent with the observations.

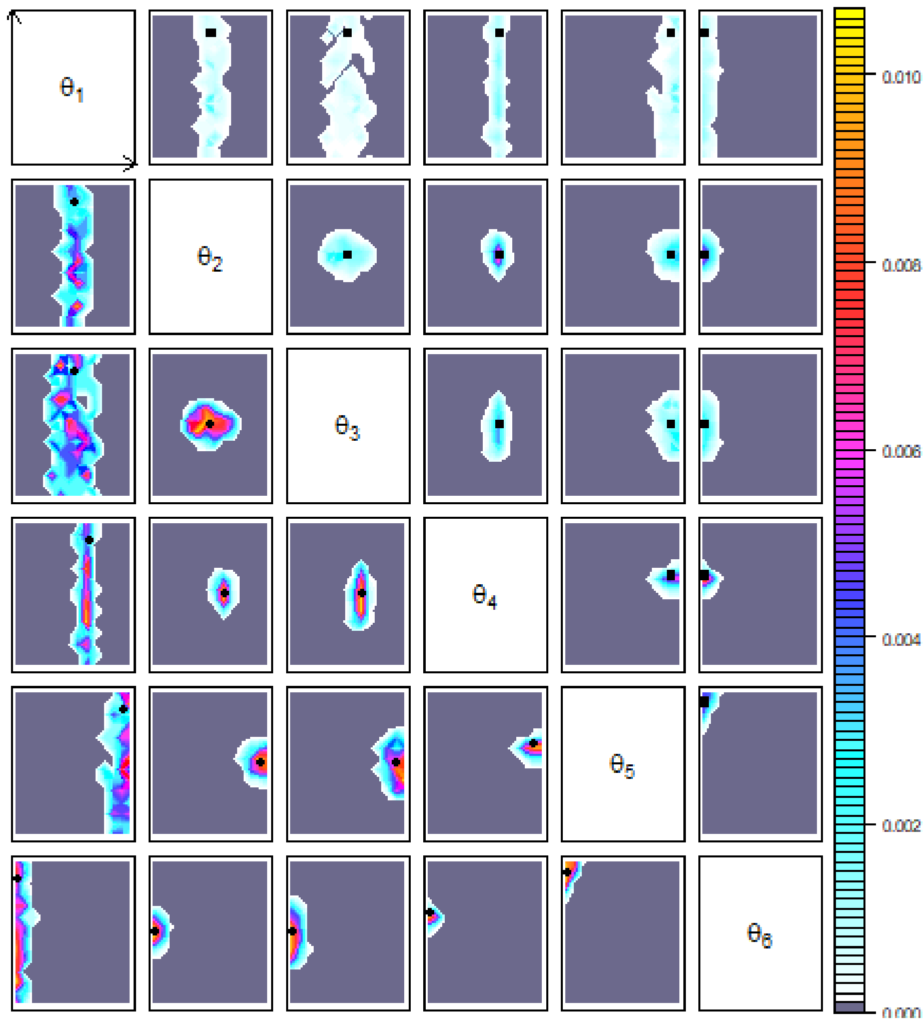

Figure S2: Density plot of the true NROY space (i.e. no emulation involved), for each pair of parameters. For each cell in a particular pairwise plot, we average across the remaining 4 parameters, and plot the proportion of these runs that are in NROY space. The axes are reversed for the lower left plots, with the colour scale set individually for each plot. The black point corresponds to $\boldsymbol{\theta}^*$.

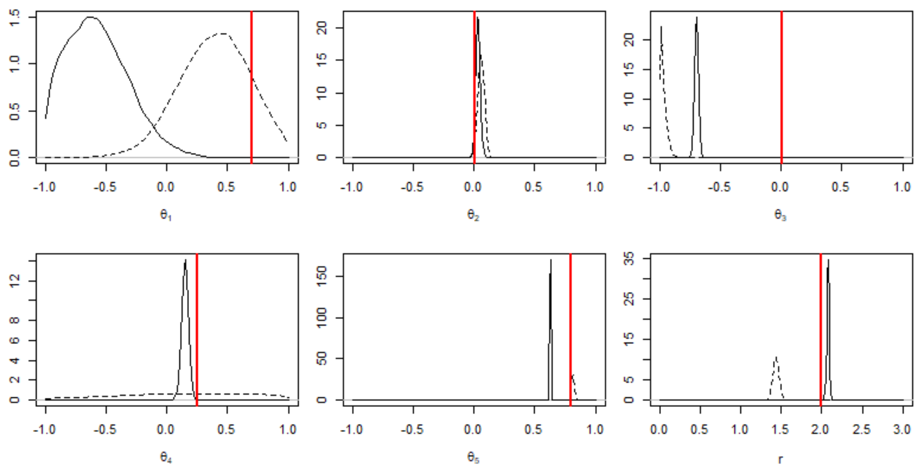

Figure S3: The posterior distributions for $\theta_1, \ldots, \theta_5$ and r, when we calibrate probabilistically on the field with the SVD basis derived from the wave 1 ensemble, $\mathbf{\Gamma}_4$ (dotted lines), and the wave 1 rotated basis (solid lines). The red vertical lines show the location of $\boldsymbol{\theta}^*$.

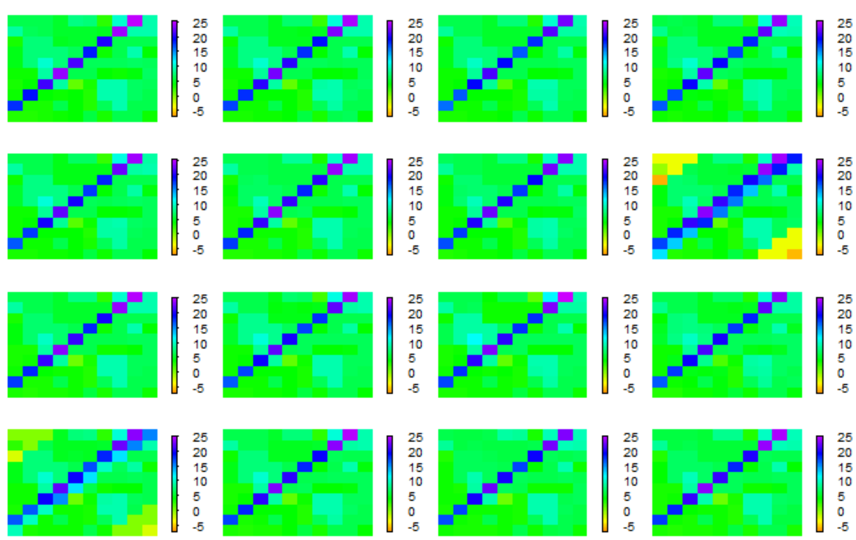

Figure S4: $f(\boldsymbol{\theta})$ at 16 samples of $\boldsymbol{\theta}$ from the calibration posterior distribution, when we emulate and calibrate with the truncated SVD basis $\mathbf{\Gamma}_4$.

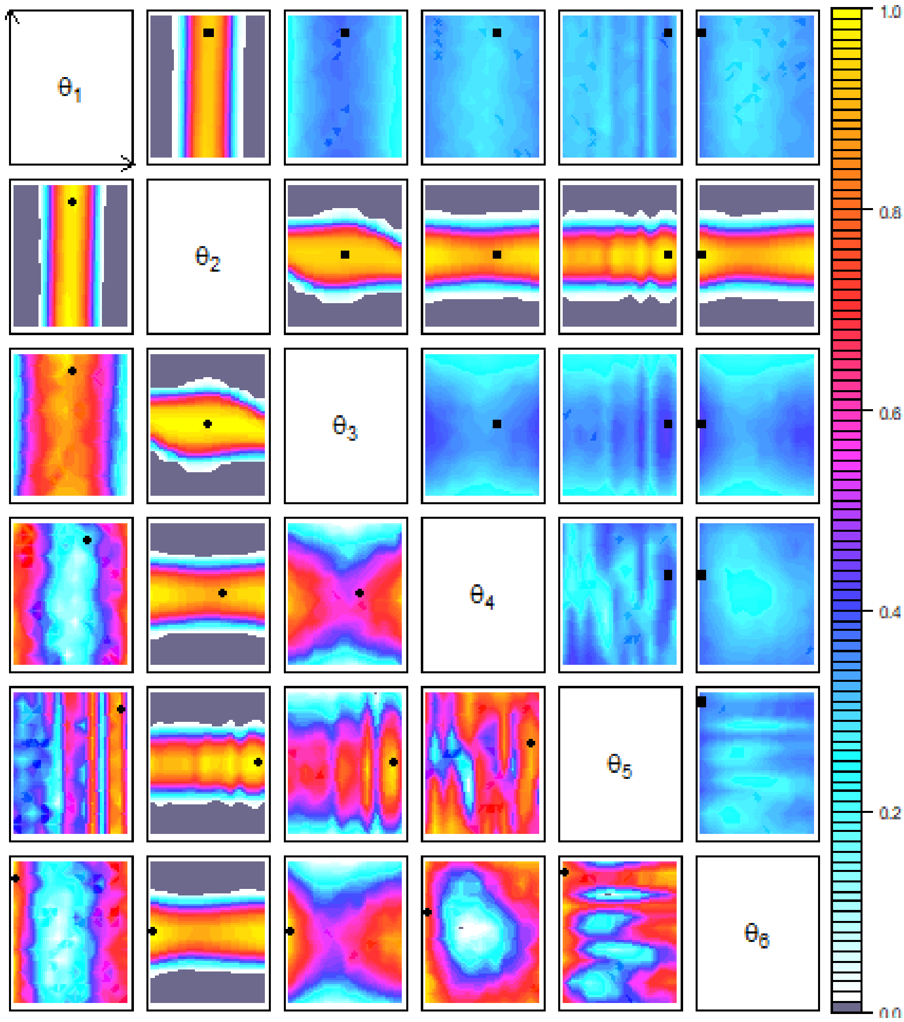

Figure S5: Density plot for the wave 1 NROY space given by history matching using the rotated basis, for each pair of parameters. As in Figure S2, for each plot we average over the remaining parameters and plot the proportion in NROY space for each pair. The axes are reversed for the lower left plots, and the colour scales set differently. The black point corresponds to $\boldsymbol{\theta}^*$.

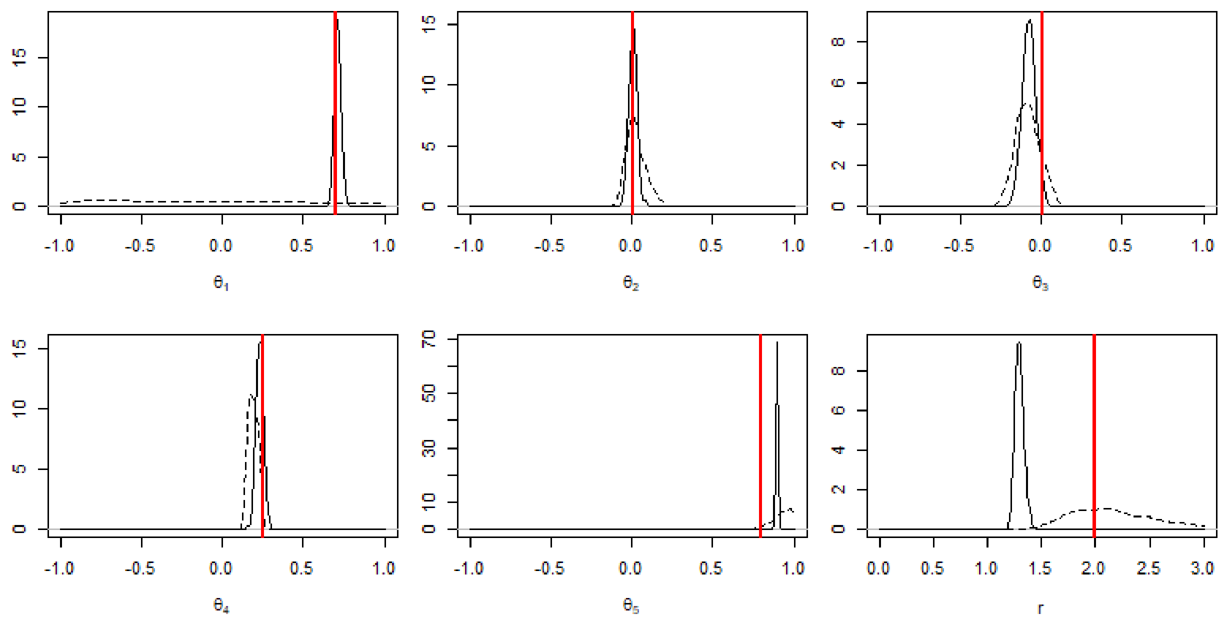

Figure S6: The posterior distributions for $\theta_1, \ldots, \theta_5$ and r, at wave 2 (solid lines) and wave 3 (dotted lines), with the red lines equal to $\boldsymbol{\theta}^*$.

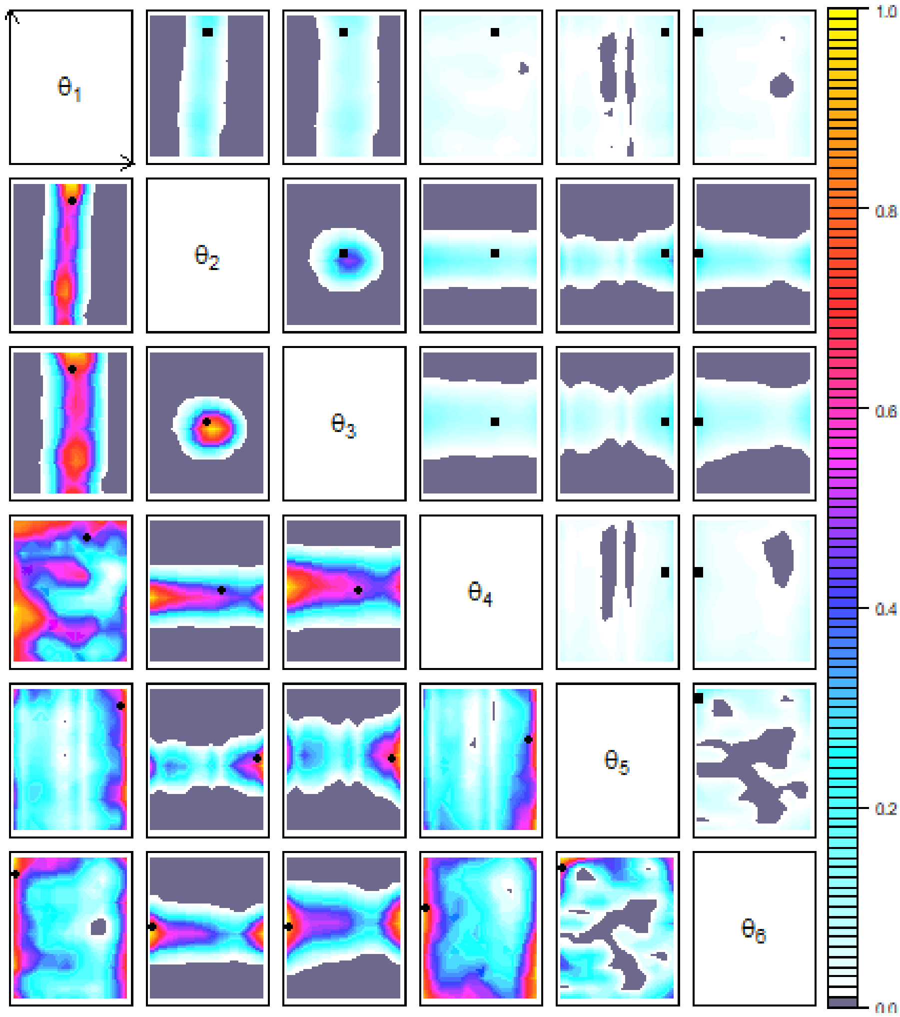

Figure S7: Density plot for the wave 2 NROY space for each pair of parameters. Regions coloured grey indicate that there are no parameter settings in NROY space here, hence we see that we have significantly reduced space, compared to Figure S5.

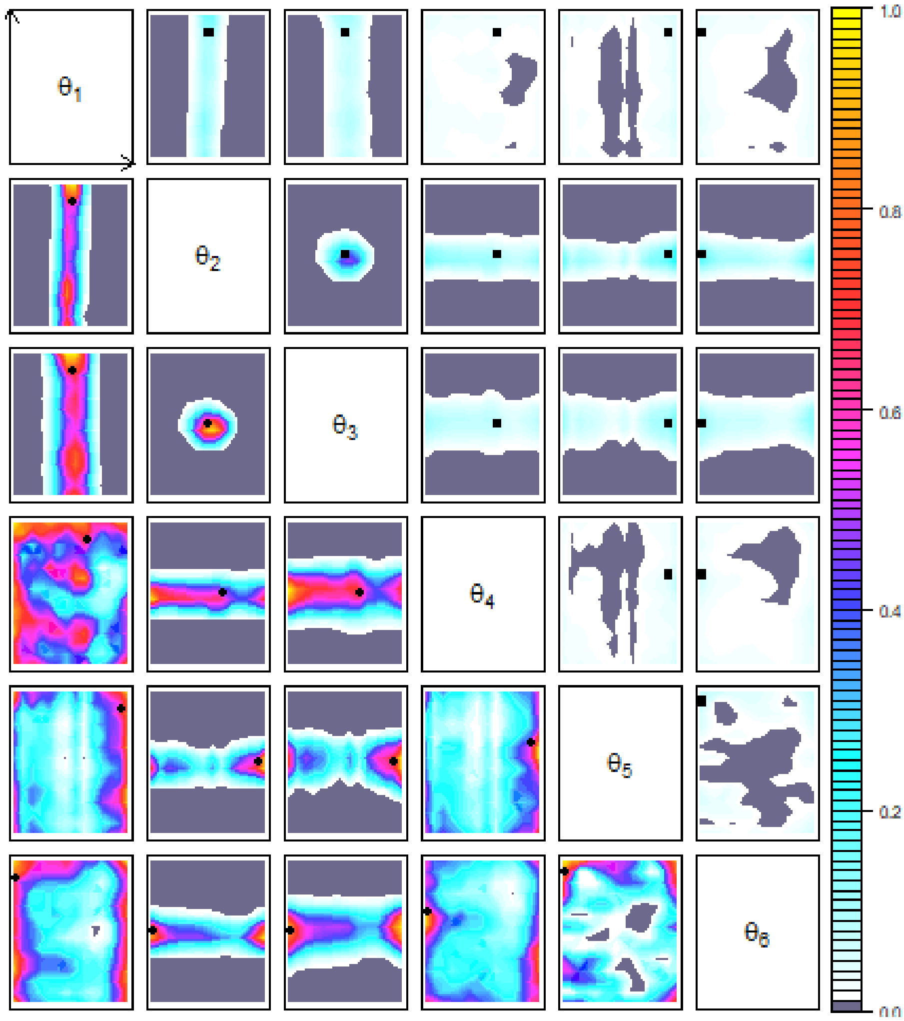

Figure S8: Density plot for the wave 3 NROY space for each pair of parameters, with space again reduced from the previous wave (we have gone from 3.1% to 2% of the original parameter space).

S2 Proofs

Weighted projection

Define $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$, a basis with dimension $\ell \times n$, and $\mathbf{W}$ an $\ell \times \ell$ positive definite matrix. The reconstruction $\mathbf{r} = (r_1, \ldots, r_\ell)^T$ of $\mathbf{f} = (f_1, \ldots, f_\ell)^T$ has reconstruction error

$$\|\mathbf{f} - \mathbf{r}\|_{\mathbf{W}} = (\mathbf{f} - \mathbf{r})^T \mathbf{W}^{-1} (\mathbf{f} - \mathbf{r}) \tag{S4}$$

in the $\mathbf{W}$ norm, where $\mathbf{r}$ is given by coefficients $\mathbf{c} = (c_1, \ldots, c_n)^T$ with

$$\mathbf{r} = \sum_{k=1}^{n} \mathbf{b}_k c_k = \mathbf{B}\mathbf{c}$$

The vector $\mathbf{c}$ that minimises the reconstruction error with respect to the $\mathbf{W}$ norm is found by first writing (S4) in terms of $\mathbf{c}$:

$$\begin{aligned}
(\mathbf{f} - \mathbf{r})^T \mathbf{W}^{-1} (\mathbf{f} - \mathbf{r}) &= (\mathbf{f} - \mathbf{B}\mathbf{c})^T \mathbf{W}^{-1} (\mathbf{f} - \mathbf{B}\mathbf{c}) \\
&= (\mathbf{f}^T - \mathbf{c}^T \mathbf{B}^T) \mathbf{W}^{-1} (\mathbf{f} - \mathbf{B}\mathbf{c}) \\
&= \mathbf{f}^T \mathbf{W}^{-1} \mathbf{f} - \mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{f} - \mathbf{f}^T \mathbf{W}^{-1} \mathbf{B}\mathbf{c} + \mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B}\mathbf{c}
\end{aligned}$$

A scalar can be differentiated by a vector $\boldsymbol{\theta}$, with symmetric matrix $\mathbf{A}$, via

$$\begin{aligned}
\frac{\partial}{\partial \boldsymbol{\theta}} \boldsymbol{\theta}^T \mathbf{y} = \frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{y}^T \boldsymbol{\theta} = \mathbf{y}^T \\
\frac{\partial}{\partial \boldsymbol{\theta}} \boldsymbol{\theta}^T \mathbf{A} \boldsymbol{\theta} = 2\boldsymbol{\theta}^T \mathbf{A}
\end{aligned}$$

Differentiating the reconstruction error with respect to $\mathbf{c}$, with symmetric $\mathbf{W}^{-1}$, we have

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{c}} \|\mathbf{f} - \mathbf{r}\|_{\mathbf{W}} &= 0 - (\mathbf{B}^T \mathbf{W}^{-1} \mathbf{f})^T - \mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} + 2\mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} \\
&= -\mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} - \mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} + \mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} + \mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} \\
&= -2\mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} + 2\mathbf{c}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B}
\end{aligned}$$

Setting equal to zero, and solving for $\mathbf{c}$, we have

$$\begin{aligned}
0 &= -2\mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} + 2\hat{\mathbf{c}}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} \\
\implies & \hat{\mathbf{c}}^T \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} = \mathbf{f}^T \mathbf{W}^{-1} \mathbf{B} \\
\implies & \mathbf{B}^T \mathbf{W}^{-1} \mathbf{B} \hat{\mathbf{c}} = \mathbf{B}^T \mathbf{W}^{-1} \mathbf{f} \\
\implies & \hat{\mathbf{c}} = (\mathbf{B}^T \mathbf{W}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{W}^{-1} \mathbf{f}
\end{aligned}$$

With $\mathbf{W} = \mathbb{I}_\ell$, this is the usual projection equation:

$$\hat{\mathbf{c}} = (\mathbf{B}^T \mathbb{I}_\ell^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbb{I}_\ell^{-1} \mathbf{f} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{f}$$

Rotational invariance

Result 1 (Invariance of $\mathcal{R}_\mathbf{W}(\cdot, \cdot)$ to rotation). *For a rotation matrix $\boldsymbol{\Lambda}$ of dimension $k \times k$, and a set of basis vectors $\boldsymbol{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$, we have*

$$\mathcal{R}_\mathbf{W}(\mathbf{B}_k, \mathbf{z}) = \mathcal{R}_\mathbf{W}(\mathbf{B}_k \boldsymbol{\Lambda}, \mathbf{z}), \quad k = 1, \ldots, n.$$

Proof. We can rewrite the reconstruction error of the rotated basis, $\mathbf{B}_k \boldsymbol{\Lambda}$, as

$$\begin{aligned} \mathcal{R}_\mathbf{W}(\mathbf{B}_k \boldsymbol{\Lambda}, \mathbf{z}) &= \|\mathbf{z} - \mathbf{B}_k \boldsymbol{\Lambda}((\mathbf{B}_k \boldsymbol{\Lambda})^T \mathbf{W}^{-1} \mathbf{B}_k \boldsymbol{\Lambda})^{-1} (\mathbf{B}_k \boldsymbol{\Lambda})^T \mathbf{W}^{-1} \mathbf{z}\|_\mathbf{W} \\ &= \|\mathbf{z} - \mathbf{B}_k \boldsymbol{\Lambda}(\boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{z}\|_\mathbf{W} \end{aligned}$$

Rotation matrix $\boldsymbol{\Lambda}$ is invertible, with $\boldsymbol{\Lambda}^T = \boldsymbol{\Lambda}^{-1}$, by definition. We apply the identity $(\mathbf{CD})^{-1} = \mathbf{D}^{-1}\mathbf{C}^{-1}$, where $\mathbf{C}, \mathbf{D}$ are $k \times k$ invertible matrices, for $\mathbf{C} = \boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k$ and $\mathbf{D} = \boldsymbol{\Lambda}$:

$$\begin{aligned} &= \|\mathbf{z} - \mathbf{B}_k \boldsymbol{\Lambda}\boldsymbol{\Lambda}^{-1}(\boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k)^{-1} \boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{z}\|_\mathbf{W} \\ &= \|\mathbf{z} - \mathbf{B}_k (\mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k)^{-1} (\boldsymbol{\Lambda}^T)^{-1} \boldsymbol{\Lambda}^T \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{z}\|_\mathbf{W} \\ &= \|\mathbf{z} - \mathbf{B}_k (\mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k)^{-1} \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{z}\|_\mathbf{W} \\ &= \mathcal{R}_\mathbf{W}(\mathbf{B}_k, \mathbf{z}) \end{aligned}$$

where in the second line, the inverse identity has been applied a second time with $\mathbf{C} = \boldsymbol{\Lambda}^T$ and $\mathbf{D} = \mathbf{B}_k^T \mathbf{W}^{-1} \mathbf{B}_k$, giving the final result. □

S2.1 Proof of Theorem 1

Before proving Theorem 1, we first prove the following results:

Result S1 (Orthogonality of the residual basis)**.** *The residual basis, $\boldsymbol{B}_\epsilon$, calculated from $\boldsymbol{F}_{\boldsymbol{\mu}}$ and $\boldsymbol{B}_p$, is orthogonal to $\boldsymbol{B}_p$ (with respect to the $\boldsymbol{W}$ norm).*

Proof. First, we show the orthogonality (in $\mathbf{W}$) of the columns of the residual ensemble, $\mathbf{F}_{\boldsymbol{\epsilon}}$, and the columns of $\mathbf{B}_p$:

$$\begin{aligned} \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\epsilon}} &= \mathbf{B}_p^T \mathbf{W}^{-1} (\mathbf{F}_{\boldsymbol{\mu}} - \mathbf{B}_p (\mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{B}_p)^{-1} \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}}) \\ &= \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}} - (\mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{B}_p)(\mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{B}_p)^{-1} \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}} \\ &= \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}} - \mathbf{B}_p^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}} \\ &= \mathbf{0} \end{aligned} \tag{S5}$$

This zero matrix has dimension $p \times n$, i.e. the basis vectors in $\mathbf{B}_p$ are orthogonal with the vectors of $\mathbf{F}_{\boldsymbol{\epsilon}}$, with respect to the $\mathbf{W}$ norm. Using this, we obtain the result by considering the (generalised) singular value decomposition of $\mathbf{F}_{\boldsymbol{\epsilon}}^T$:

$$\begin{aligned} \mathbf{F}_{\boldsymbol{\epsilon}}^T &= \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T \\ \implies \mathbf{F}_{\boldsymbol{\epsilon}}^T &= \mathbf{U}\boldsymbol{\Sigma}\mathbf{B}_{\boldsymbol{\epsilon}}^T \end{aligned}$$

where $\mathbf{U}$ is an orthonormal $n \times n$ matrix, $\boldsymbol{\Sigma}$ is a diagonal $n \times n$ matrix, and $\mathbf{V} = \mathbf{B}_{\boldsymbol{\epsilon}}$ is an $\ell \times n$ matrix with $\mathbf{B}_{\boldsymbol{\epsilon}}^T\mathbf{W}^{-1}\mathbf{B}_{\boldsymbol{\epsilon}} = \mathbb{I}_n$. From here, we have that

$$\begin{aligned} \implies \mathbf{F}_{\boldsymbol{\epsilon}} &= \mathbf{B}_{\boldsymbol{\epsilon}}\boldsymbol{\Sigma}^T\mathbf{U}^T \\ \implies \mathbf{F}_{\boldsymbol{\epsilon}}\mathbf{U} &= \mathbf{B}_{\boldsymbol{\epsilon}}\boldsymbol{\Sigma}^T\mathbf{U}^T\mathbf{U} \\ \implies \mathbf{B}_p^T\mathbf{W}^{-1}\mathbf{F}_{\boldsymbol{\epsilon}}\mathbf{U} &= \mathbf{B}_p^T\mathbf{W}^{-1}\mathbf{B}_{\boldsymbol{\epsilon}}\boldsymbol{\Sigma}^T \end{aligned}$$

where we have multiplied on the left by $\mathbf{B}_p^T\mathbf{W}^{-1}$. From (S5), we have $\mathbf{B}_p^T\mathbf{W}^{-1}\mathbf{F}_{\boldsymbol{\epsilon}} = \mathbf{0}$, hence

$$\implies \mathbf{B}_p^T\mathbf{W}^{-1}\mathbf{B}_{\boldsymbol{\epsilon}}\boldsymbol{\Sigma}^T = \mathbf{0}$$

The final $p+1$ eigenvalues on the diagonal of $\boldsymbol{\Sigma}$ are zero (the p vectors in $\mathbf{B}_p$, and the ensemble mean, remove $p+1$ degrees of freedom). Therefore, we are only interested in the leading $n-p-1$ columns of $\mathbf{B}_{\boldsymbol{\epsilon}}$, as all of the variability in $\mathbf{F}_{\boldsymbol{\mu}}$ has already been explained. By discarding the columns associated with zero eigenvalues, we have that

$$\implies \mathbf{B}_p^T\mathbf{W}^{-1}[\mathbf{B}_{\boldsymbol{\epsilon}}]_{n-p-1} = \mathbf{0}$$

because $\boldsymbol{\Sigma}$ is diagonal, and hence we have the result. □

From Result (S1), we can show that the basis vector selected at step k of the optimal rotation algorithm, $\boldsymbol{\gamma}_k^* = \mathbf{B}_\epsilon^{k-1}\boldsymbol{\lambda}_k$, is orthogonal to those previously selected, $\boldsymbol{\Gamma}_{k-1}^* = (\boldsymbol{\gamma}_1^*, \dots, \boldsymbol{\gamma}_{k-1}^*)$:

$$(\boldsymbol{\Gamma}_{k-1}^*)^T\mathbf{W}^{-1}\boldsymbol{\gamma}_k^* = (\boldsymbol{\Gamma}_{k-1}^*)^T\mathbf{W}^{-1}\mathbf{B}_\epsilon^{k-1}\boldsymbol{\lambda}_k = \mathbf{0}.$$

Result S2. *When $\boldsymbol{B}$ is a basis for $\boldsymbol{F}_{\boldsymbol{\mu}}$, we can write $\boldsymbol{B}_\epsilon^k = \boldsymbol{B}\boldsymbol{\Lambda}_\epsilon^k$ for square $\boldsymbol{\Lambda}_\epsilon^k$, i.e. the residual basis at iteration k of the algorithm contains linear combinations of the vectors of $\boldsymbol{B}$, and hence each vector selected by the algorithm is a linear combination of $\boldsymbol{B}$.*

Proof. By the singular value decomposition of the residual ensemble after selecting k basis vectors, we have

$$(\mathbf{F}_\epsilon^k)^T = \mathbf{U}_\epsilon^k\boldsymbol{\Sigma}_\epsilon^k(\mathbf{B}_\epsilon^k)^T$$

for orthonormal $\mathbf{U}_\epsilon^k$, and diagonal $\mathbf{\Sigma}_\epsilon^k$. We can write $\mathbf{F}_{\boldsymbol{\mu}} = \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}$ ($\mathbf{B}$ is a basis for $\mathbf{F}_{\boldsymbol{\mu}}$, hence the ensemble is a linear combination of the basis vectors), for $n \times n$ matrix $\mathbf{\Lambda}_{\boldsymbol{\mu}}$. At iteration k of the optimal rotation algorithm, we have

$$\begin{aligned}\mathbf{B}_\epsilon^k \mathbf{\Sigma}_\epsilon^k (\mathbf{U}_\epsilon^k)^T &= \mathbf{F}_{\boldsymbol{\mu}} - \mathbf{\Gamma}_k^* ((\mathbf{\Gamma}_k^*)^T \mathbf{W}^{-1} \mathbf{\Gamma}_k^*)^{-1} (\mathbf{\Gamma}_k^*)^T \mathbf{W}^{-1} \mathbf{F}_{\boldsymbol{\mu}} \\ \implies \mathbf{B}_\epsilon^k &= (\mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}} - \mathbf{\Gamma}_k^* ((\mathbf{\Gamma}_k^*)^T \mathbf{W}^{-1} \mathbf{\Gamma}_k^*)^{-1} (\mathbf{\Gamma}_k^*)^T \mathbf{W}^{-1} \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}) \mathbf{U}_\epsilon^k (\mathbf{\Sigma}_\epsilon^k)^{-1}\end{aligned}$$

Set $k = 1$, i.e. we have only selected one basis vector so far. When $k = 1$, the algorithm selects a linear combination of $\mathbf{B}$ by construction, hence we have $\boldsymbol{\gamma}_1^* = \mathbf{\Gamma}_1^* = \mathbf{B}\tilde{\boldsymbol{\lambda}}_1$ for vector $\tilde{\boldsymbol{\lambda}}_1$. Therefore,

$$\begin{aligned}\mathbf{B}_\epsilon^1 &= (\mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}} - \mathbf{B}\tilde{\boldsymbol{\lambda}}_1 ((\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^{-1} (\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}) \mathbf{U}_\epsilon^1 (\mathbf{\Sigma}_\epsilon^1)^{-1} \\ &= \mathbf{B}(\mathbf{\Lambda}_{\boldsymbol{\mu}} - \tilde{\boldsymbol{\lambda}}_1 ((\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^{-1} (\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}) \mathbf{U}_\epsilon^1 (\mathbf{\Sigma}_\epsilon^1)^{-1} \\ &= \mathbf{B}\mathbf{\Lambda}_\epsilon^1\end{aligned} \tag{S6}$$

with $\mathbf{\Lambda}_\epsilon^1 = \mathbf{\Lambda}_{\boldsymbol{\mu}} - \tilde{\boldsymbol{\lambda}}_1 ((\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^{-1} (\mathbf{B}\tilde{\boldsymbol{\lambda}}_1)^T \mathbf{W}^{-1} \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}) \mathbf{U}_\epsilon^1 (\mathbf{\Sigma}_\epsilon^1)^{-1}$. Then at iteration $k = 2$, we optimise over linear combinations of $\mathbf{B}_\epsilon^1$, so that

$$\boldsymbol{\gamma}_2^* = \mathbf{B}_\epsilon^1 \boldsymbol{\lambda}_2 = \mathbf{B}\mathbf{\Lambda}_\epsilon^1 \boldsymbol{\lambda}_2 = \mathbf{B}\tilde{\boldsymbol{\lambda}}_2$$

i.e. the second basis vector is a linear combination of $\mathbf{B}$. It follows that at iteration k, we select a new basis vector where

$$\boldsymbol{\gamma}_k^* = \mathbf{B}_\epsilon^{k-1} \boldsymbol{\lambda}_k = \mathbf{B}\mathbf{\Lambda}_\epsilon^{k-1} \boldsymbol{\lambda}_k = \mathbf{B}\tilde{\boldsymbol{\lambda}}_k$$

for $\mathbf{\Lambda}_\epsilon^{k-1} = \mathbf{\Lambda}_{\boldsymbol{\mu}} - \tilde{\mathbf{\Lambda}}_{k-1} ((\mathbf{B}\tilde{\mathbf{\Lambda}}_{k-1})^T \mathbf{W}^{-1} \mathbf{B}\tilde{\mathbf{\Lambda}}_{k-1})^{-1} (\mathbf{B}\tilde{\mathbf{\Lambda}}_{k-1})^T \mathbf{W}^{-1} \mathbf{B}\mathbf{\Lambda}_{\boldsymbol{\mu}}) \mathbf{U}_\epsilon^1 (\mathbf{\Sigma}_\epsilon^1)^{-1}$, and $\tilde{\mathbf{\Lambda}}_{k-1} = (\tilde{\boldsymbol{\lambda}}_1, \ldots, \tilde{\boldsymbol{\lambda}}_{k-1})$ (so that $\mathbf{\Gamma}_{k-1}^* = \mathbf{B}\tilde{\mathbf{\Lambda}}_{k-1}$). Therefore, we have that the residual basis at each iteration is a linear combination of the original basis, $\mathbf{B}$, and hence each new basis vector is. □

Theorem 1. $\mathbf{\Gamma}^*$ *in step 3 of the optimal rotation algorithm is an orthogonal rotation of* $\boldsymbol{B}$*.*

Proof. Assume that we have performed k iterations of the algorithm, resulting in the basis

$$\mathbf{\Gamma}^* = (\boldsymbol{\gamma}_1^*, \ldots, \boldsymbol{\gamma}_k^*, [\mathbf{B}_\epsilon^k]_{n-k}) \tag{S7}$$

With $\mathbf{B}_\epsilon^k = \mathbf{B}\mathbf{\Lambda}_\epsilon^k$ and $\boldsymbol{\gamma}_j^* = \mathbf{B}\tilde{\boldsymbol{\lambda}}_j$ (Result S2), we rewrite (S7) as

$$\mathbf{\Gamma}^* = (\mathbf{B}\tilde{\boldsymbol{\lambda}}_1, \ldots, \mathbf{B}\tilde{\boldsymbol{\lambda}}_k, [\mathbf{B}\mathbf{\Lambda}_\epsilon^k]_{n-k}) = \mathbf{B}(\tilde{\boldsymbol{\lambda}}_1, \ldots, \tilde{\boldsymbol{\lambda}}_k, [\mathbf{\Lambda}_\epsilon^k]_{n-k}) = \mathbf{B}\mathbf{\Lambda}.$$

We show that $\mathbf{\Lambda}^T\mathbf{\Lambda} = \mathbb{I}_n$, i.e. $\mathbf{\Lambda}$ is a rotation matrix. We have

$$\mathbf{\Lambda}^T\mathbf{\Lambda} = \begin{pmatrix} \tilde{\boldsymbol{\lambda}}_1^T\tilde{\boldsymbol{\lambda}}_1 & \tilde{\boldsymbol{\lambda}}_1^T\tilde{\boldsymbol{\lambda}}_2 & \dots & \tilde{\boldsymbol{\lambda}}_1^T[\mathbf{\Lambda}_\epsilon^k]_{n-k} \\ \tilde{\boldsymbol{\lambda}}_2^T\tilde{\boldsymbol{\lambda}}_1 & \tilde{\boldsymbol{\lambda}}_2^T\tilde{\boldsymbol{\lambda}}_2 & \dots & \tilde{\boldsymbol{\lambda}}_2^T[\mathbf{\Lambda}_\epsilon^k]_{n-k} \\ \vdots & & \ddots & \\ [\mathbf{\Lambda}_\epsilon^k]_{n-k}^T\tilde{\boldsymbol{\lambda}}_1 & [\mathbf{\Lambda}_\epsilon^k]_{n-k}^T\tilde{\boldsymbol{\lambda}}_2 & \dots & [\mathbf{\Lambda}_\epsilon^k]_{n-k}^T[\mathbf{\Lambda}_\epsilon^k]_{n-k} \end{pmatrix} \tag{S8}$$

The upper-left $k \times k$ block can be written as

$$\tilde{\boldsymbol{\lambda}}_i^T\tilde{\boldsymbol{\lambda}}_j = \tilde{\boldsymbol{\lambda}}_i^T\mathbf{B}^T\mathbf{W}^{-1}\mathbf{B}\tilde{\boldsymbol{\lambda}}_j = (\boldsymbol{\gamma}_i^*)^T\mathbf{W}^{-1}\boldsymbol{\gamma}_j^* = \begin{cases} 1 & \text{if} \quad i = j \\ 0 & \quad \text{otherwise.} \end{cases}$$

Similarly,

$$[\mathbf{\Lambda}_\epsilon^k]_{n-k}^T\tilde{\boldsymbol{\lambda}}_j = (\mathbf{B}[\mathbf{\Lambda}_\epsilon^k]_{n-k})^T\mathbf{W}^{-1}\boldsymbol{\gamma}_j^* = \mathbf{B}_\epsilon^T\mathbf{W}^{-1}\boldsymbol{\gamma}_j^* = \mathbf{0},$$

by Result S1. Finally,

$$[\mathbf{\Lambda}_\epsilon^k]_{n-k}^T[\mathbf{\Lambda}_\epsilon^k]_{n-k} = [\mathbf{\Lambda}_\epsilon^k]_{n-k}^T\mathbf{B}^T\mathbf{W}^{-1}\mathbf{B}[\mathbf{\Lambda}_\epsilon^k]_{n-k} = [\mathbf{B}_\epsilon^k]_{n-k}^T\mathbf{W}^{-1}[\mathbf{B}_\epsilon^k]_{n-k} = \mathbb{I}_{n-k}$$

and hence from (S8) we have $\mathbf{\Lambda}^T\mathbf{\Lambda} = \mathbb{I}_n$, and $\mathbf{\Lambda}$ is a rotation matrix. □

S2.2 Gram-Schmidt invariance

Gram-Schmidt orthonormalisation imposes orthonormality on basis vectors $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ (Björck, 1967). It can be written in terms of matrices (Björck, 1994):

$$\mathbf{B} = \mathbf{\Gamma}\mathbf{R}$$

where $\mathbf{\Gamma}$ is an $l \times n$ basis containing normalised, orthogonal vectors $\boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_n$, and $\mathbf{R}$ is an $n \times n$ upper-triangular matrix. Therefore, the j^{th} new basis vector is a linear combination of the first j basis vectors of $\mathbf{B}$.

Result S3 (Gram-Schmidt invariance)**.** *The reconstruction given by the first q vectors of the original basis is equal to the reconstruction given by the first q vectors of the orthogonal basis:*

$$\mathcal{R}_{\boldsymbol{W}}(\boldsymbol{B}_q, \boldsymbol{z}) = \mathcal{R}_{\boldsymbol{W}}(\mathbf{\Gamma}_q, \boldsymbol{z}), \quad q = 1, \dots, n$$

Proof. Using $\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{\Gamma} = \mathbb{I}_n$ (i.e. we've imposed orthonormality with respect to the $\mathbf{W}$ norm), we have

$$\begin{aligned}\mathbf{B}(\mathbf{B}^T\mathbf{W}^{-1}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{W}^{-1}\mathbf{z} &= \mathbf{\Gamma}\mathbf{R}((\mathbf{\Gamma}\mathbf{R})^T\mathbf{W}^{-1}\mathbf{\Gamma}\mathbf{R})^{-1}(\mathbf{\Gamma}\mathbf{R})^T\mathbf{W}^{-1}\mathbf{z}\\ &= \mathbf{\Gamma}\mathbf{R}(\mathbf{R}^T\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{\Gamma}\mathbf{R})^{-1}\mathbf{R}^T\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{z}\\ &= \mathbf{\Gamma}\mathbf{R}(\mathbf{R}^T\mathbf{R})^{-1}\mathbf{R}^T\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{z}\end{aligned}$$

As $\mathbf{R}$ is invertible, and applying the identity $(\mathbf{C}\mathbf{D})^{-1} = \mathbf{D}^{-1}\mathbf{C}^{-1}$ for square matrices $\mathbf{C}$ and $\mathbf{D}$, we have

$$\begin{aligned}&= \mathbf{\Gamma}\mathbf{R}\mathbf{R}^{-1}\mathbf{R}^{-T}\mathbf{R}^T\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{z}\\ &= \mathbf{\Gamma}\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{z}\\ &= \mathbf{\Gamma}(\mathbf{\Gamma}^T\mathbf{W}^{-1}\mathbf{\Gamma})^{-1}\mathbf{\Gamma}^T\mathbf{z}\end{aligned}$$

i.e. the reconstruction of $\mathbf{z}$ is the same with both $\mathbf{B}$ and $\mathbf{\Gamma}$. The proof proceeds analogously for any truncation of these bases. □

Alternatively, this result could be shown by proving that both $\mathbf{B}$ and $\mathbf{\Gamma}$ span the same subspace, and using that a basis gives unique representations of general fields (Kuttler, 2012).

S3 CanAM4

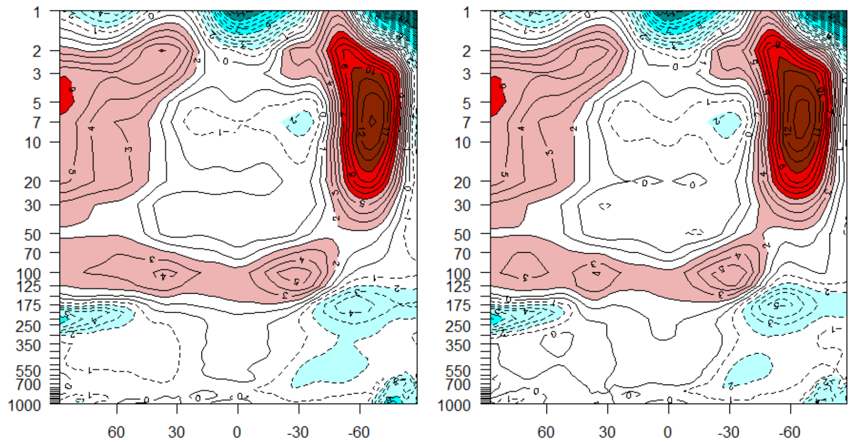

Figure S9: The TA anomaly for the standard run (left), and for run 005 of the new ensemble, the 'best' run in the wave 2 ensemble, in terms of minimising the root mean squared error. The large warm bias from the standard run has not been removed.

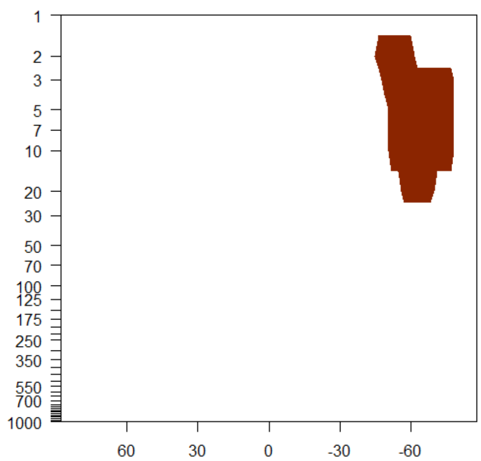

Figure S10: The grid boxes where we deem there to be a potential structural error for TA, and hence where we increase the discrepancy, as described in Section 5.

References

Björck, Å. (1967). Solving linear least squares problems by Gram-Schmidt orthogonalization. *BIT Numerical Mathematics 7*(1), 1–21.

Björck, Å. (1994). Numerics of Gram-Schmidt orthogonalization. *Linear Algebra and Its Applications 197*, 297–316.

Kuttler, K. (2012). *Linear algebra: theory and applications.* The Saylor Foundation.