Editors: Isabel Beichl, isabel.beichl@nist.gov
Julian V. Noble, jvn@virginia.edu

FITTING NATURE'S BASIC FUNCTIONS PART IV: THE VARIABLE PROJECTION ALGORITHM

By Bert W. Rust

IN THE PREVIOUS INSTALLMENTS, WE DESCRIBED LINEAR AND NONLINEAR LEAST SQUARES CALCULATIONS. THE LATTER ARE CONSIDERABLY MORE DIFFICULT THAN THE FORMER. WE WILL NOW EXAMINE THE VARIABLE PROJECTION ALGORITHM, WHICH OFTEN GREATLY simplifies nonlinear least squares calculations because it does not require iteration on parameters that appear linearly in the model.

Nonlinear Least Squares

In Part III of this series,[1] we used nonlinear least squares to fit the exponential model

$$y(t-t_0) = C_0 e^{\beta(t-t_0)} \ , \tag{1}$$

where t_0 = 1856.0, to the annual global total fossil fuel production record that Gregg Marland and his colleagues compiled.[2] Since then, they have added another year of data to the data set, and the production figures for the years 1982 to 1998 have been slightly revised. You can get the updated record at http://cdiac.ornl.gov/trends/emis/em_cont.htm.

Fitting Equation 1 to the new data gives the parameter estimates

$$\hat{C}_0 = 158 \pm 11 \ \ [\text{Mt C}] , \tag{2}$$

$$\hat{\beta} = 0.02697 \pm .00053 \ \ [\text{yr}^{-1}] . \tag{3}$$

The fit, which Figure 1 plots, accounts for 97.5 percent of the total variance in the data, but the residuals, plotted in Figure 2, exhibit a quasicyclic variation whose amplitude appears to be growing along with the exponential baseline.

The residuals' Fourier power spectrum is similar to the one in Figure 4 of Part II[3] for the residuals from the reduced quadratic fit to the global temperature record. Here, the dominant peak is at ≈65.6 years; there, it was at ≈61.4 years. The simplest assumption about the amplitude's growth rate is that it is increasing at the same exponential rate as the baseline. This suggests a model of the form

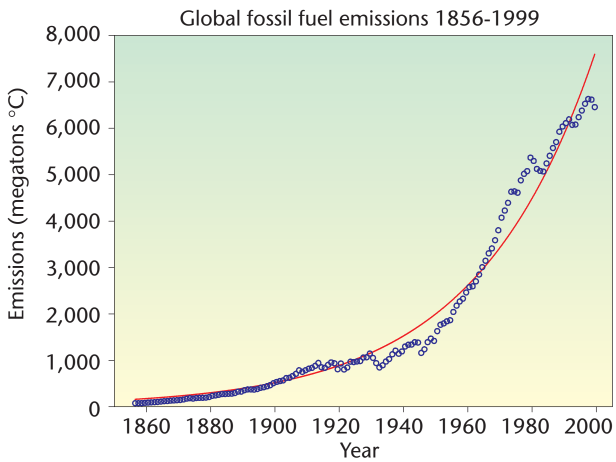

Figure 1. The nonlinear fit of the simple exponential model in Equation 1 to the global total fossil fuel emissions for the years 1856 through 1999.

$$y(t-t_0) = c_1 e^{\alpha_1(t-t_0)} + A e^{\alpha_1(t-t_0)} \sin\left[\frac{2\pi}{\alpha_2}(t-t_0+\phi)\right], \tag{4}$$

with unknown parameters c_1, α_1, A, α_2, and ϕ, which would require nonlinear least squares for their estimation. The work done so far provides initial estimates $c_1^{(0)} = 160$, $\alpha_1^{(0)} = 0.027$, and $\alpha_2^{(0)} = 65$, but getting good initial estimates $A^{(0)}$ and $\phi^{(0)}$ would require some effort.

In Part III, we recommended rewriting models involving sinusoids like the one in Equation 4 by defining new parameters

$$c_2 = A\cos\left(\frac{2\pi\phi}{\alpha_2}\right), \qquad c_3 = A\sin\left(\frac{2\pi\phi}{\alpha_2}\right), \tag{5}$$

so that we can write the model

$$y(t-t_0) = c_1 e^{\alpha_1(t-t_0)} + c_2 e^{\alpha_1(t-t_0)} \sin\left[\frac{2\pi(t-t_0)}{\alpha_2}\right] + c_3 e^{\alpha_1(t-t_0)} \cos\left[\frac{2\pi(t-t_0)}{\alpha_2}\right], \tag{6}$$

with the parameters A and ϕ being replaced by c_2 and c_3. We can get estimates of A and ϕ from estimates of c_2, c_3, and α_2 by

$$\hat{A} = +\sqrt{\hat{c}_2^2 + \hat{c}_3^2}, \qquad \hat{\phi} = \frac{\hat{\alpha}_2}{2\pi} \tan^{-1}\left(\frac{\hat{c}_3}{\hat{c}_2}\right), \tag{7}$$

and we can compute uncertainties in those estimates via Equations 40 and 41 in Part III.[1]

Of course, finding initial estimates $c_2^{(0)}$ and $c_3^{(0)}$ is no easier than finding $A^{(0)}$ and $\phi^{(0)}$, but Equation 6 expresses the model in a form that depends linearly on three of the unknown parameters and nonlinearly on only two of them. This makes it possible to get the fit via a *variable projection algorithm*,[4] which iterates only on the nonlinear parameters and computes the linear parameters by linear least squares—with no need for initial estimates.

The Variable Projection Method

The variable projection method assumes that we can write the fitting function as

$$y(t;\mathbf{c},\alpha) = \sum_{j=1}^{n} c_j \Phi_j(t;\alpha) \tag{8}$$

where

$$\mathbf{c} \equiv (c_1, c_2, \ldots, c_n)^T \tag{9}$$

is an n-vector of linear parameters,

$$\alpha \equiv (\alpha_1, \alpha_2, \ldots, \alpha_p)^T \tag{10}$$

is a p-vector of nonlinear parameters, and the $\Phi_j(t,\alpha)$ are functions that depend only on the nonlinear parameters. The notation is not meant to imply that all the Φ_j's depend on all the α's. Any given Φ_j can depend on all, some, or none of the α's, but none of them depend on any of the c_j's.

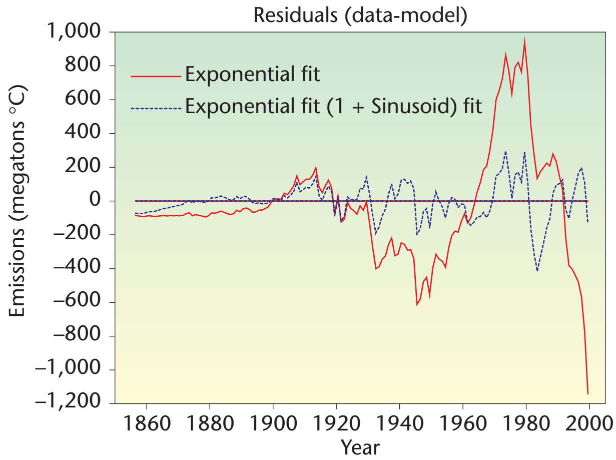

Figure 2. The residuals for exponential fits to the fossil fuel emissions record. The solid curve gives the residuals for the fit shown in Figure 1; the dashed line gives the residuals for the fit of Equation 6 shown in Figure 3.

Given m observations (t_i, y_i), we define

$$\mathbf{y} \equiv (y_1, y_2, \ldots, y_m)^T, \tag{11}$$

and, for any given α,

$$\Phi(\alpha) \equiv \begin{bmatrix} \Phi_1(t_1;\alpha) & \Phi_2(t_1;\alpha) & \ldots & \Phi_n(t_1;\alpha) \\ \Phi_1(t_2;\alpha) & \Phi_2(t_2;\alpha) & \ldots & \Phi_n(t_2;\alpha) \\ \cdot & \cdot & & \cdot \\ \cdot & \cdot & & \cdot \\ \cdot & \cdot & & \cdot \\ \Phi_1(t_m;\alpha) & \Phi_2(t_m;\alpha) & \ldots & \Phi_n(t_m;\alpha) \end{bmatrix}, \tag{12}$$

then we can write the sum of squared residuals that must be minimized as

$$\mathcal{L}(\mathbf{c},\alpha) = \sum_{i=1}^{m} \left[y_i - \sum_{j=1}^{n} c_j \Phi_j(t_i;\alpha) \right]^2 \tag{13}$$

$$= \left[\mathbf{y} - \Phi(\alpha)\mathbf{c}\right]^T \left[\mathbf{y} - \Phi(\alpha)\mathbf{c}\right] \tag{14}$$

$$\equiv \left\| \mathbf{y} - \Phi(\alpha)\mathbf{c} \right\|^2. \tag{15}$$

This objective function must be minimized with respect to the elements of both $\mathbf{c}$ and α, but it is advantageous to write the minimization problem as

$$\mathcal{L}(\hat{\mathbf{c}},\hat{\alpha}) = \min_{\alpha}\left\{\min_{\mathbf{c}}\left\{\left\|\mathbf{y}-\Phi(\alpha)\mathbf{c}\right\|^2\right\}\right\} \tag{16}$$

and note that for any fixed α, the inner minimization is a linear least square problem just like the one in Part I.[5] Therefore, for any given α, the inner minimum occurs at

$$\mathbf{c}_{\min}(\alpha) = \Phi^{\dagger}(\alpha)\mathbf{y} , \tag{17}$$

where the $n \times m$ matrix

$$\Phi^{\dagger}(\alpha) \equiv \left[\Phi^T(\alpha)\Phi(\alpha)\right]^{-1}\Phi^T(\alpha) \tag{18}$$

is called the *generalized inverse* or *pseudoinverse* of the $m \times n$ matrix $\Phi(\alpha)$. This means that we can write the minimization problem in Equation 16 as

$$\mathcal{L}(\hat{\mathbf{c}},\hat{\alpha}) = \min_{\alpha}\left\{\left\|\mathbf{y}-\Phi(\alpha)\Phi^{\dagger}(\alpha)\mathbf{y}\right\|^2\right\}, \tag{19}$$

which depends only on the parameters α.

Because the objective function is very nonlinear, we need an iteration to isolate the minimizer $\hat{\alpha}$, but once we find it, we can calculate the estimate $\hat{\mathbf{c}}$ from

$$\hat{\mathbf{c}} = \left[\Phi^T(\hat{\alpha})\Phi(\hat{\alpha})\right]^{-1}\Phi^T(\hat{\alpha})\mathbf{y} \tag{20}$$

with no iteration being required.

Gene Golub and Victor Pereya[4] first gave an iterative algorithm for solving the minimization problem in Equation 19; they also provided a Fortran computer code called VARPRO, which combined it with a linear least squares algorithm to solve for $\hat{\mathbf{c}}$. To fully appreciate the magnitude of their accomplishment, recall that minimization algorithms require, at each step, the derivatives of the objective function with respect to the unknown parameters. How would you like to try to differentiate the objective function in Equation 19 with respect to the vector α? The advantage of a variable separable algorithm is so apparent to real-world modelers that it is amazing that VARPRO isn't used more widely. You can get VARPRO online at www.netlib.org/opt/varpro. A good review detailing its development and applications over the last 30 years will soon appear in the journal *Inverse Problems*.[6]

Back to the Fossil Fuel Emissions

Equation 6 is already written in the VARPRO form (Equation 8) with $p = 2$, $n = 3$, and

$$\Phi_1(t-t_0;\alpha) = e^{\alpha_1(t-t_0)} \tag{21}$$

$$\Phi_2(t-t_0;\alpha) = e^{\alpha_1(t-t_0)}\sin\left[\frac{2\pi(t-t_0)}{\alpha_2}\right] \tag{22}$$

$$\Phi_3(t-t_0;\alpha) = e^{\alpha_1(t-t_0)}\cos\left[\frac{2\pi(t-t_0)}{\alpha_2}\right]. \tag{23}$$

VARPRO requires the partial derivatives of these three functions with respect to α_1 and α_2 in order to compute the partial derivatives of Equation 19's objective function. Using initial estimates $\alpha_1^{(0)} = 0.027$ and $\alpha_2^{(0)} = 65$ leads, after 16 iterations, to the estimates

$$\begin{aligned}
&\hat{c}_1 = 133.8 \pm 5.0 \ \ [MtC],\\
&\hat{\alpha}_1 = 0.02814 \pm .00034 \ \ [\mathrm{yr}^{-1}],\\
&\hat{c}_2 = -21.5 \pm 4.4 \ \ [MtC],\\
&\hat{c}_3 = 13.5 \pm 5.0 \ \ [MtC],\\
&\hat{\alpha}_2 = 64.9 \pm 1.5 \ \ [\mathrm{yr}].
\end{aligned} \tag{24}$$

The amplitude and phase estimates corresponding to $\hat{c}_2$ and $\hat{c}_3$ are

$$\begin{aligned}
&\hat{A} = 25.4 \pm 1.6 \ \ [MtC]\\
&\hat{\phi} = 26.6 \pm 2.7 \ \ [\mathrm{yr}].
\end{aligned} \tag{25}$$

The relatively small uncertainty in $\hat{A}$ seems to indicate that the added oscillation is statistically significant, but using the F-test described in Part II is required to ensure that the addition of the three new parameters did produce a statistically significant reduction in the residual variance.[3] The null hypothesis is

$$H_0: \quad c_2^* = 0, \qquad c_3^* = 0, \qquad \alpha_2^* = 0, \tag{26}$$

so $k = 3$ and $n = 5$, and the number of measurements was $m = 144$. The sum of squared residuals for the simple exponential fit was $SSR_H = 1.4546 \times 10^7$ and for the exponential-sinusoidal model, $SSR_F = 1.575 \times 10^6$. Therefore, we have

$$u = \frac{\mathrm{SSR}_H - \mathrm{SSR}_F}{\mathrm{SSR}_F} \cdot \frac{m-n}{k}$$
$$= 381.6 >> F_{0.95}(3{,}139) = 2.670 \quad , \tag{27}$$

so the null hypothesis is resoundingly rejected.

The new fit, which Figure 3 plots, accounts for 99.7 percent of the total variance in the data. The residuals, plotted as a dashed line in Figure 2, show a marked improvement over those for the simple exponential fit, but they do not look like white noise. The largest single systematic change is in the years 1979 through 1983, which was the period of the Iranian revolution and the Iran–Iraq war. This drop in emissions shows up clearly in the plots of the emissions data, but the corresponding relative drop was only about 5.6 percent, which is much smaller than the 16.6 percent drop in the years 1943 through 1945 and the 26.0 percent drop in the years 1929 through 1932. You could easily accommodate these one-time blips by assuming linear drops during the relevant time periods, but three additional linear parameters would be required in the model.

Global Temperatures Again

In Part III, we fit a model[1] of the form

$$\phi(t,\alpha) = \alpha_1 + \alpha_2 \exp\left[\alpha_3(t-t_0)\right] + \alpha_5 \sin\left[\frac{2\pi(t-t_0)}{\alpha_4}\right] + \alpha_6 \cos\left[\frac{2\pi(t-t_0)}{\alpha_4}\right], \tag{28}$$

using six nonlinear parameters $\alpha_1, \ldots, \alpha_6$, to the Climatic Research Unit's annual global average temperature anomaly record (www.cru.uea.ac.uk/cru/cru.htm). That fit would have been easier using VARPRO because we could also write the model as

$$y(t-t_0) = c_1 + c_2 e^{\alpha_1(t-t_0)} + c_3 \sin\left[\frac{2\pi(t-t_0)}{\alpha_2}\right] + c_4 \cos\left[\frac{2\pi(t-t_0)}{\alpha_2}\right], \tag{29}$$

which has only two nonlinear parameters α_1 and α_2. These two parameters correspond exactly to the α_1 and α_2 in Equation 6 for the fossil fuel emissions, meaning α_1 is an exponential rate constant and α_2 is the period of a sinusoid. If we define

$$A = +\sqrt{c_3^2 + c_4^2}\ , \qquad \phi = \frac{\alpha_2}{2\pi}\tan^{-1}\left(\frac{c_4}{c_3}\right), \tag{30}$$

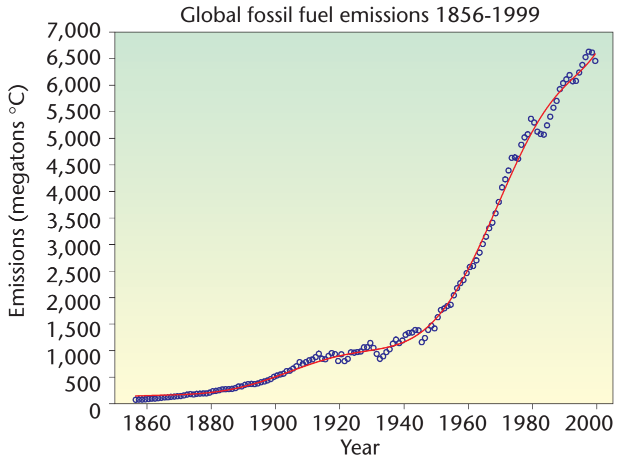

Figure 3. Nonlinear fit of the exponential–sinusoidal model in Equation 6 to the global total fossil fuel emissions for the years 1856 through 1999.

then we can also write Equation 29 as

$$y(t-t_0) = c_1 + c_2 e^{\alpha_1(t-t_0)} + A\sin\left[\frac{2\pi}{\alpha_2}(t-t_0+\phi)\right], \tag{31}$$

which corresponds to the form in Equation 4 of the fossil fuel emissions model.

Table 1 compares the estimates of α_1, α_2, and ϕ for the two models. The values in the second column are from Equations 24 and 25 and the discussion following them, and the values in the third column are from the fifth column of Table 2 in Part III,[1] although the names of the parameters differ. Comparing the estimates for the parameters yields the following interesting observations:

- Halving the $\hat{\alpha}_1$ for the fossil fuel emissions model gives a value that lies well inside the $\pm 1\sigma$ interval for the corresponding global temperature estimate.
- The two $\hat{\alpha}_2$ estimates are the same.
- The difference between the two $\hat{\phi}$ estimates is approximately one half of the common estimate of $\hat{\alpha}_2$.

Table 1. Parameter estimates and statistics for fossil fuel emissions and global average temperature anomalies.

Parameter	Fossil fuel emissions	Global temperatures
$\hat{\alpha}_1$ [yr^{-1}]	0.02814 ± .00034	0.0160 ± .0028
$\hat{\alpha}_2$ [yr]	64.9 ± 1.5	64.9 ± 2.1
$\hat{\phi}$ [yr]	26.6 ± 2.7	–5.1 ± 2.4
SSR	1.575 × 106	1.476
100 R2	99.73 percent	80.98 percent

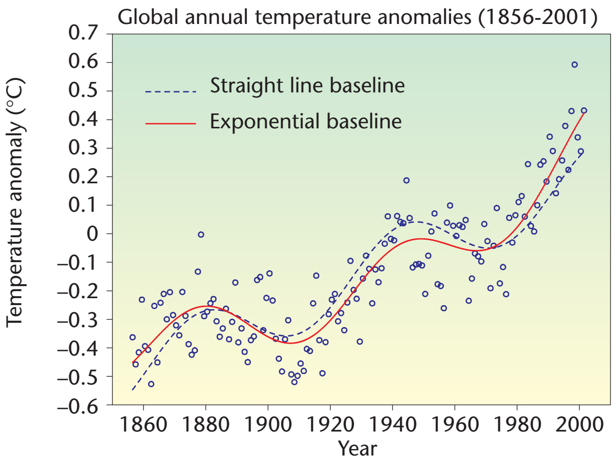

Figure 4. Linear fits of the exponential plus sinusoid model in Equation 33 and the straight line plus sinusoid model in Equation 34 to the global annual temperature anomalies for the years 1856 through 2001.

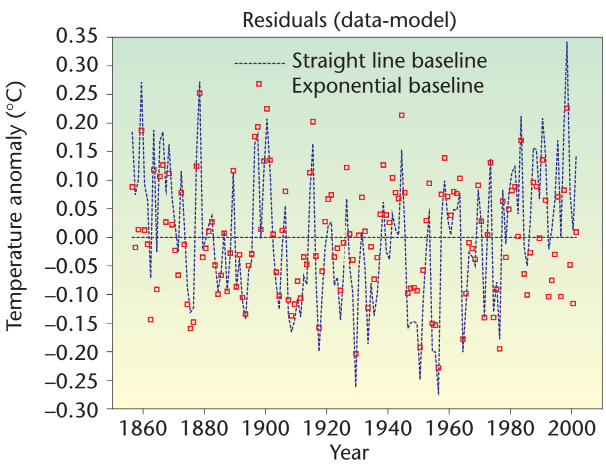

Figure 5. Residuals for the two fits plotted in Figure 4.

Although the first point could be just a coincidence, it is highly unlikely that the second two are the result of chance. For more than two cycles, the same oscillation occurs in both records with a phase shift between them of almost exactly one half of a cycle. I worked with Bernadette Kirk[7] in 1982 on this inverse correlation of the variations in the two records; we suggested that it might represent a feedback mechanism by which rising temperatures suppress increases in the production of fossil fuels. We also suggested a dynamical model

$$\frac{dP}{dt} = \left(\alpha - \beta\frac{dT}{dt}\right)P\,, \qquad P(t_0) = P_0 \tag{32}$$

relating the production $P(t)$ and temperature $T(t)$, with the parameters α, β, and P_0 to be determined by fitting. Frank Crosby and I extended this analysis in 1994;[8] we allowed for a lag in the response to temperature changes and added one-time innovations to model World War II and the Great Depression.

Regardless of whether temperature changes exert a Gaiaen feedback on fossil fuel production, the parameter estimates in Table1 suggest a simplified model

$$\begin{aligned} y(t-t_0) &= c_1 + c_2 \exp[0.01407(t-t_0)] \\ &\quad + c_3 \sin\left[\frac{2\pi}{64.9}(t - t_0 - 5.85)\right] \end{aligned} \tag{33}$$

for the temperature anomaly record. The exponential rate constant is one half of the estimate for the fossil fuel production record, and I adjusted the phase shift to differ by exactly one half cycle from the phase shift estimate for that record.

Fitting this model to the data is an easy linear least squares problem. The result is plotted as a solid curve in Figure 4, and the residuals are plotted as discrete points in Figure 5. For comparison, we also fitted models with linear and reduced quadratic baselines,

$$\begin{aligned} y(t-t_0) &= c_1 + c_2 \begin{Bmatrix} (t-t_0) \\ (t-t_0)^2 \end{Bmatrix} \\ &\quad + c_3 \sin\left[\frac{2\pi}{64.9}(t - t_0 - 5.85)\right]. \end{aligned} \tag{34}$$

Table 2 gives the parameter estimates and some statistics for all three fits. The linear baseline fit is displayed as a dashed line in Figure 4, but to avoid confusion, the reduced quadratic baseline fit is not plotted because it tracks the exponential baseline fit so closely. The residuals for the linear baseline are plotted as a dashed curve in Figure 5. Comparing them with the residuals for the exponential baseline fit clearly reveals that the data require a curved baseline. Because the linear baseline model accounts for only 73.92 percent of the record's variance (whereas the reduced quadratic and exponential baseline models account for 80.71 percent and 80.81 percent, respectively), we can safely conclude not only that the Earth's troposphere is warming but also that this warming is accelerating. Because of the close agreement between the fits for the reduced quadratic and exponential baseline models, determining whether that acceleration is constant or something worse is not possible.

In Part II, we discovered that adding a linear term to the reduced quadratic baseline model did not produce a statis-

Table 2. Parameter estimates and statistics for fits of Equations 33, 34, and 35 to the global average temperature anomalies.

Parameter	Straight line	Reduced quadratic	Exponential	Full quadratic
$\hat{c}_1$	$-0.494 \pm .020$	$-0.380 \pm .013$	$-0.510 \pm .017$	$-0.355 \pm .026$
$\hat{c}_2$	$(4.74 \pm .24) \times 10^{-3}$	$(3.27 \pm .14) \times 10^{-5}$	$0.1094 \pm .0045$	$(3.85 \pm .54) \times 10^{-5}$
$\hat{c}_3$	$0.112 \pm .015$	$0.105 \pm .013$	$0.104 \pm .013$	$0.103 \pm .013$
$\hat{c}_4$				$(-9.0 \pm 8.1) \times 10^{-4}$
SSR	2.024359	1.497461	1.490011	1.484507
100 R2	73.92 percent	80.71 percent	80.81 percent	80.88 percent

tically significant reduction in the sum of squared residuals.[3] This would seem to imply that the data require a monotone-increasing baseline, but that analysis did not take the 64.9-year sinusoidal variation into account. We now repeat that test by fitting the model

$$y(t-t_0) = c_1 + c_4(t-t_0) + c_2(t-t_0)^2 + c_3 \sin\left[\frac{2\pi}{64.9}(t-t_0-5.85)\right] \tag{35}$$

to the data. That fit's parameter estimates and statistics appear in the last column of Table 2. The high relative uncertainty in the estimate $\hat{c}_4$ indicates that it is not statistically significant, and a formal F-test confirms this. For the null hypothesis, we take

$$H_0: \qquad c_4^* = 0, \tag{36}$$

so

$$\begin{aligned} u &= \frac{\mathrm{SSR}_H - \mathrm{SSR}_F}{\mathrm{SSR}_F} \cdot \frac{m-n}{k} \\ &= \frac{1.497461 - 1.484507}{1.484507} \cdot \frac{146-4}{1} \\ &= 1.2391 < F_{0.95}(1,142) = 3.9077 \end{aligned} \tag{37}$$

and the null hypothesis is accepted.

The residuals plotted in Figure 5 for the exponential baseline model are certainly more random than those for the straight-line baseline model, but they still exhibit systematic periodic variations that are not easily discernible by inspecting the plot. We will address those variations in a later installment, but even with the present incomplete model in Equation 33, the residual standard error is

$$\sigma_r = \sqrt{\frac{1}{m-n}\mathrm{SSR}} = 0.1021\ [^\circ\mathrm{C}], \tag{38}$$

and the net increase predicted by the model for the interval 1856 through 2001 is 0.8740°C.

If we provisionally accept that fit as the signal and the residuals as the noise, then the global warming in the last 146 years is more than 8.5 times greater than the noise level. But the analysis is not complete until we find a model that reduces the residuals to white noise. Future installments will discuss diagnostic tests to determine when that result has been achieved.

References

1. B.W. Rust, "Fitting Nature's Basic Functions Part III: Exponentials, Sinusoids, and Nonlinear Least Squares," *Computing in Science & Eng.*, vol. 4, no. 4, 2002, pp. 72–77.
2. G. Marland, T.A. Boden, and R.J. Andres, "Global, Regional, and National CO_2 Emissions," *Trends: A Compendium of Data on Global Change*, Oak Ridge Nat'l Lab., 2000.
3. B.W. Rust, "Fitting Nature's Basic Functions Part II: Estimating Uncertainties and Testing Hypotheses," *Computing in Science & Eng.*, vol. 3, no. 6, 2001, pp. 60–64.
4. G.H. Golub, and V. Pereyra, "The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Whose Variables Separate," *SIAM J. Numerical Analysis*, vol. 10, no. 2, 1973, pp. 413–432.
5. B.W. Rust, "Fitting Nature's Basic Functions Part I: Polynomials and Linear Least Squares," *Computing in Science & Eng.*, vol. 3, no. 5, 2001, pp. 84–89.
6. G. Golub and V. Pereyra, "Separable Nonlinear Least Squares: The Variable Projection Method and Its Applications," submitted to *Inverse Problems*, 2003.
7. B.W. Rust and B.L. Kirk, "Modulation of Fossil Fuel Production by Global Temperature Variations," *Environment International*, vol. 7, 1982, pp. 419–422.
8. B.W. Rust and F.J. Crosby, "Further Studies on the Modulation of Fossil Fuel Production by Global Temperature Variations," *Environment International*, vol. 20, no. 4, 1994, pp. 429–456.

Bert W. Rust is a research mathematician at the National Institute of Standards and Technology. His research interests include ill-posed problems, time-series modeling, nonlinear regression, and observational cosmology. He received a BS in engineering physics and an MS in mathematics from the University of Tennessee and a PhD in astronomy from the University of Illinois. Contact him at NIST, 100 Bureau Dr. Stop 8910, Gaithersburg, MD 20899-8910; bwr@cam.nist.gov.