7. Errors and Uncertainty Estimation

The estimation of the *FIRAS* errors and uncertainties was a major part of the *FIRAS* data reduction, calibration and destriping. This chapter is an attempt to distill the output of that effort into a concise yet complete description of what is known about the *FIRAS* errors and uncertainties.

Uncertainties in the *FIRAS* data come from many different sources and manifest themselves in many different ways. A full covariance matrix is $210 * 6067 \times 210 * 6067$ which is too large to be convenient. Fortunately many of the uncertainties are quite well described by a few dominant terms that show the source of the uncertainty.

We concentrate on the HIGH, LOWF and HRES data sets as these are the best. The individual channels and scan modes have many of the same uncertainty estimates but post-destriper data sets have lower systematic errors.

7.0.1. Sensitivity

The *FIRAS* sensitivity is characterized by the C and D vectors, which are estimates for the random errors. Over most of the sky and most of the frequency range there is a positive signal. However, there are systematic errors of several types. The differences are important, as some of the systematic errors can be reduced by averaging in some ways but not in others. The uncertainties are both frequency and position dependent. The systematic uncertainties are frequency and position dependent in different ways.

7.0.2. Formal Calibration Uncertainty

In the calibration program (Section 5 and Appendix C) the derivatives of the results with respect to the calibration parameters are calculated. Although the derivatives are expressly calculated for the calibration data the same derivatives can be used for the sky data to propagate the uncertainties of the sky data and the calibration model to the calibrated sky. For convenience the derivatives and the propagation are divided into groups: the detector measurements, the calibration emissivities, the bolometer model parameters, the temperature measurements of all but the XCAL, and the temperature measurement of the XCAL. In earlier discussions (Fixsen *et al.* 1994b), this covariance matrix was expressed:

$$V = D + P_0 E P_0 + JCJ + PUP + PTP \tag{21}$$

D is due to random uncertainties in the readings on the detectors. P_0EP_0 is due to uncertainties in the emissivities derived in the calibration. JCJ is due to uncertainties in the other parameters derived from the calibration. PUP is due to random errors in reading the thermometers while observing the sky. PTP is due to uncertainty in the absolute temperature scale of the external calibrator.

The description is complicated by the destriping process, but it is still useful to follow these terms. At the end of this chapter we add sections on tests for other errors and a description of how these uncertainty estimates can be used in practical problems.

7.1. The Detector Noise

Conceptually the detector noise is simple. At each readout there is some noise on the detector which then gets processed and ends up in the final output. The detector outputs are contaminated by noise from the resistance of the detectors, the thermal conductance, cosmic rays, and the readout noise of the electronics. The Johnson noise, and intrinsic noise on the thermal link can be well approximated by a white noise source followed by a filter. The electronic readout noise is small and can also be well approximated by a white noise source followed by a (different) filter. The cosmic rays can be approximated by delta functions applied at random times with some distribution. Although all of these sources lead to noise that is correlated in the time domain, because the noise processes are *stationary* they are uncorrelated in the frequency domain. Thus the only frequency correlations are introduced by the apodization.

7.1.1. The D Matrix

There are two ways to estimate the matrices. As the coadds are formed from the averages of a number of IFGs, the mean dispersion is calculated and from it, the covariance in the final data. This matrix is termed the "D Matrix". For historical reasons we write the covariance as D_ν^2/N_p where D^2 is a diagonal matrix and N_p is some number of "observations" at each pixel, even though the matrix is not really diagonal. By keeping only the square root of the diagonal part we store it in a vector space ("D Vector"). As such we have estimates for the uncertainty of an IFG in any scan mode. For weighted averages we use a standard weight and estimate a "D Vector" "per observation".

This is done for the calibration and the sky data. The calibration data taken in the sky-like near null condition with all controllables near 2.7 K show the same noise characteristics as the sky data. The calibration data taken with one or more of the controllables at a high temperature had higher noise, so these data were deweighted with respect to the sky-like null calibration data. But, the higher signals meant in general they had higher signal to noise ratios (Section 5.5).

The data near the Galactic plane also had higher noise because the data had significant variation over a small change in position. These variations were fit before coadding by a second template to reduce their impact on the data processing. Data that had many glitches removed had more noise that those that had fewer glitches removed so the weights were adjusted to account for the number of glitches (Section 4.4). Although the glitch noise has a redder spectrum than the other bolometer noise we used a single spectrum to allow a simple weight for each IFG, coadd and pixel.

7.1.2. The C Matrix

As the data are destriped, coadds from different times and different channels are compared. The *rms* dispersions of the coadds are accumulated to find the mean covariance matrix (the "C" matrix). In the destriper each coadd is assigned a weight depending on the channel, scan mode, and glitch rate. A single noise spectrum is then determined from the dispersions among the coadds. The noise spectra of the different channels are different in shape as well as amplitude so the use of a single noise spectrum for the combined data is not as efficient as it could have been. However the use of a frequency dependent weight would only reduce the uncertainties $\sim 1\%$ at the price of requiring a frequency dependent weight, and position.

The sky covariance is given by:

$$C_{\nu\nu'} = \frac{\sum_i \left[S_{i\nu}^{(sky)} - \sum_k J_{k\nu} f_{ki} - A_{p(i)\nu}\right] w_i \left[S_{i\nu'}^{(sky)} - \sum_k J_{k\nu'} f_{ki} - A_{p(i)\nu'}\right] / (1 - f_i)}{N - n_J} \quad (22)$$

where $N \equiv \sum_{i \in sky} 1$, w_i is the weight (Section 6.2), f_i is the fraction of pixel weight supplied by the ith coadd, and n_J is the number of model functions (between 10 and 30). It gives an estimate of the variance of a single interferogram that includes both the random error due to the detector plus any remaining systematic variance. It also provides the scale of the other destriper errors.

The square root of the diagonal is termed the "C vector". Because the dispersion can include errors associated with long term drift and channel to channel errors it might be expected to be larger than the "D vector". But because the weights have been adjusted for time and position the "C vector" is actually lower than the "D vector" for some frequencies. The "C vector" is more inclusive and it is the preferred uncertainty estimator. The C and D "vectors" are only a few percent different (Figure 7.1). In the following text C_ν refers to the "vector" and $C_{\nu\nu'}$ refers to the matrix.

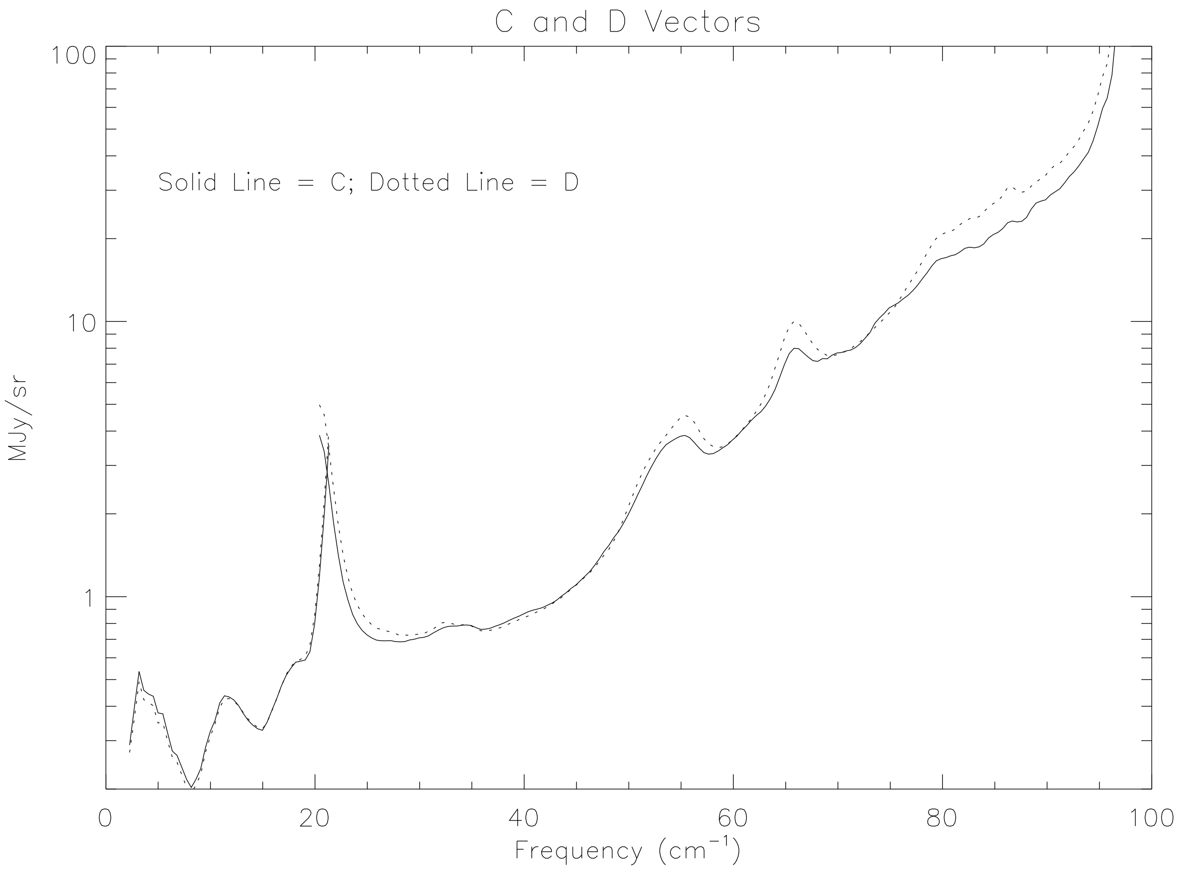

Fig. 7.1.— Relative noise verses frequency — C and D vectors for LOWF (2 to 21cm^{-1}) and HIGH (20 to 100 cm^{-1}).

The off diagonal elements are estimated from the data exactly as the diagonal parts are. This yields a 210×210 covariance matrix. It is convenient to separate it into a correlation part A and the diagonal part C which is the "C vector" identically, $C_{\nu\nu'} = C_\nu C_{\nu'} A_{\nu\nu'}$. A brief look at the $A_{\nu\nu'}$, shows it is very regular. Within 1% it is summarized as $A_{\delta\nu}$, where $\delta\nu = |\nu - \nu'|$. This is called the "A vector".

The expected apodization effect can be calculated from $A = FFT(apod^2)$. This nearly reproduces the measured "A vector" (Figure 7.2). The differences can be attributed to glitch correction effects, baseline subtraction, and perturbations like the vibrations and

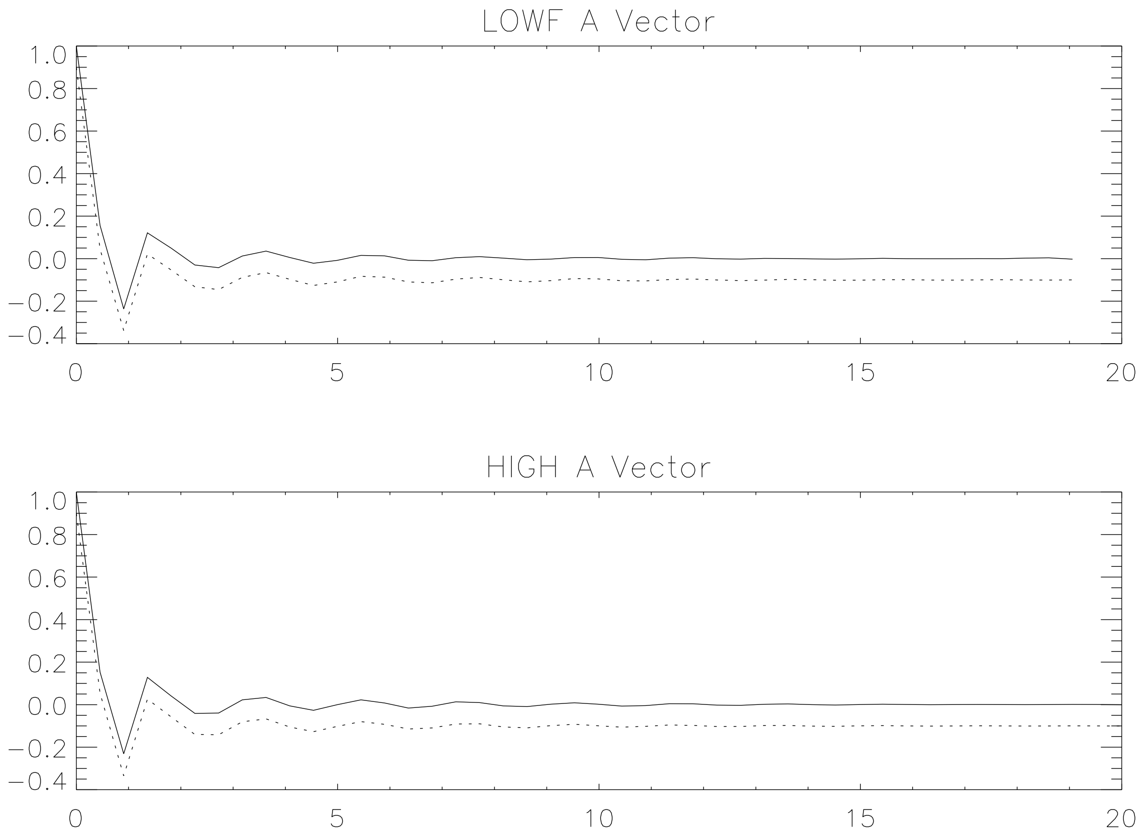

Fig. 7.2.— Frequency-frequency correlations — A vectors for LOWF and HIGH; The solid lines are the "A Vector" estimated from the correlations in the C matrix. The dotted line is the "A Vector" estimated from the apodization, and offset by 0.1 for clarity.

harmonics. The fact that final sky dispersions have the correlations implied by the apodization function alone to within $\sim 1\%$ gives strong support to this uncertainty description.

7.2. The Emissivity Uncertainty

The P_0EP_0 term is the uncertainty which results from the "emissivity" terms of the calibration model in the calibration program, FISH, (*e.g.* errors in the emissivity of the internal calibrator). These errors are propagated to the final sky data by multiplying by the Planck function, P_0, associated with the temperature of the emitting object: in the covariance matrix this term takes the form $P_0 \cdot E \cdot P_0^T$. Since the critical objects are kept at 2.73 K during the observations the P_0EP_0 is summarized with P_0 =Planck(2.73) with the emissivity uncertainty matrix E. But the P_x (P for the external calibrator) is the actual sky spectrum.

For the low frequencies ($\nu < 20$ cm^{-1} or 600 GHz) the sky spectrum is well approximated by a 2.73 K Planck spectrum and the full calculation can be done. The resulting error estimate is contained in the PEP_Offset field of the error term file. For high frequencies ($\nu > 25$ cm^{-1} or 750 GHz) the sky is quite variable and the calculation must be done at each pixel. Most of the terms of the calculation are essentially zero because Planck$\rightarrow$0 at high frequencies. Only the single term P_x, is significant and it assumes the form of a gain error. This gain error, P_gEP_g, is contained in the PEP_GAIN field of the error term file. For frequencies in the range $20 - 25$ cm^{-1} this situation is more complicated but is reasonably well modeled by using both P_0EP_0 and P_gEP_g.

Since the emissivities are calculated independently at each frequency, P_0EP_0 and P_gEP_g have no correlations across frequencies. However, because the whole sky is calibrated with the same instrument emissivity model all pixels are correlated. The terms become important whenever spectra from a large area of the sky are averaged to form a single spectrum.

7.2.1. The Destriper and P_0EP_0 Offset Errors

The uncertainty in the parameters of a linear least-squares fit is given by the diagonal elements of the inverse of the *curvature matrix* (Section 6.4). Thus, the uncertainties in the correction spectra of the destriper are given by:

$$\Delta J_{k\nu} \;=\; C_\nu \sqrt{[\mathbf{Q}]_{kk}} \tag{23}$$

These quantities are useful for determining whether a particular function of the model is significant. The off-diagonal elements of $\mathbf{Q}$ give the covariance among the different model functions.

The destriper introduces correlations between pixels through the correction spectra, $J_k(\nu)$. These correlations and the uncertainties on the pixel spectra, $A_p(\nu)$, are contained in the lower right sub-matrix $\boldsymbol{\Upsilon}$. In principle this matrix is easy to compute, however, in practice it is not easy to store. However, if we perform a Cholesky decomposition on the real symmetric matrix $\mathbf{Q}$, then we can express this matrix as:

$$\boldsymbol{\Upsilon} \;=\; \mathbf{D^{-1}} + \beta\beta^{\mathbf{T}} \tag{24}$$

where $\beta \;\equiv\; \mathbf{D^{-1}R^{T}L}$ and $\mathbf{Q} = \mathbf{L}\;\mathbf{L^{T}}$. Note that β has the same dimensions as $\mathbf{R}$, where $\mathbf{R}$ is the matrix of stripe kernels (Section 6.2).

The excluded Galactic pixels are not included in $\boldsymbol{\Upsilon}$, however, it can be shown that correlations among the pixel spectra:

$$A_{p\nu} \;=\; \frac{\sum_{i\in p}\left[S_{i\nu}^{(sky)} \;-\; \sum_k J_{k\nu} f_{ki}\right] n_i}{\sum_{i\in p} n_i} \tag{25}$$

introduced by the destriper corrections, are given by the same expression: $\mathbf{D^{-1}} + \mathbf{D^{-1}R^{T}QRD^{-1}}$, where $\mathbf{D^{-1}}$ and $\mathbf{R}$ now include all observed pixels. ($\mathbf{Q}$) still excludes the pixels from the Galactic center. The diagonal elements of $\boldsymbol{\Upsilon}$ give the pixel spectra variance in units of $C_{\nu\nu'}$.

When modeling the pixel spectra, whether the entire sky or a more localized area, it is necessary to compute the corresponding weights matrix given by $\left[\boldsymbol{\Upsilon}(p, p' \in R)\right]^{-\mathbf{1}}$ where p and p' are pixels within the region R. If the number of pixels within the region is small then this can be done directly, however it is also possible to perform the inverse so that only a small 10 × 10 matrix need be inverted explicitly. If we express the covariance matrix (or the relevant sub-matrix) as:

$$\boldsymbol{\Upsilon}(p, p' \in R) = \widetilde{\mathbf{D}}^{-\mathbf{1/2}} \;\left[\mathbf{I_m} + \widetilde{\mathbf{D}}^{\mathbf{1/2}}\tilde{\beta}\tilde{\beta}^{\mathbf{T}}\widetilde{\mathbf{D}}^{\mathbf{1/2}}\right]\widetilde{\mathbf{D}}^{-\mathbf{1/2}} \tag{26}$$

$$= \;\widetilde{\mathbf{D}}^{-\mathbf{1/2}}\;\left[\mathbf{I_m} + \mathbf{U\Phi\Phi^{T}U^{T}}\right]\;\widetilde{\mathbf{D}}^{-\mathbf{1/2}} \tag{27}$$

where $\widetilde{\mathbf{D}} \;\equiv\; \mathbf{D}(p \in R)$, $\mathbf{I_m}$ is an $m \times m$ identity matrix (m = number of pixels in region), $\tilde{\beta} \;\equiv\; \beta(p \in R)$ and $\mathbf{U}$ is an $m \times n$ row orthogonal matrix (n = number of model functions) then:

$$\mathbf{W} \equiv \left[\boldsymbol{\Upsilon}(\mathbf{p}, \mathbf{p}' \in R)\right]^{-\mathbf{1}} \;=\; \widetilde{\mathbf{D}}^{\mathbf{1/2}}\left[\mathbf{I_m} - \mathbf{U\Lambda\Lambda^{T}U^{T}}\right]\widetilde{\mathbf{D}}^{\mathbf{1/2}} \;=\; \widetilde{\mathbf{D}} - \mathbf{\Omega\Omega^{T}} \tag{28}$$

where

$$\mathbf{\Lambda\Lambda^T} \equiv \mathbf{I_n} - (\mathbf{I_n}+\mathbf{\Phi\Phi^T})^{-1}. \tag{29}$$

and

$$\mathbf{\Omega} \equiv \widetilde{\mathbf{D}}^{\mathbf{1/2}}\ \mathbf{U\Lambda} \tag{30}$$

We can compute $\mathbf{U}$ and $\mathbf{\Phi}$ through a QR decomposition of $\widetilde{D}^{1/2}\widetilde{\beta}$ and obtain $\mathbf{\Lambda}$ via a Cholesky decomposition of $\mathbf{I_n} - (\mathbf{I_n}+\mathbf{\Phi\Phi^T})^{-1}$. Note that $\mathbf{\Omega}$ is an $m \times n$ matrix and thus easy to store and manipulate.

Because the destriper model includes the difference between the calibration and Planck spectra, the destriper uncertainties, β, include "improved" P_0EP_0 and JCJ offset uncertainties. These P_0EP_0 uncertainties are most important when modeling a large portion of the sky, for in this case the error is dominated by the uncertainty of the calibration rather than that of the sky. This can be illustrated as follows. Consider a destriper model with a single entire mission offset with the sky described by a single pixel of N observations. If the total number of calibration observations is given by n, then:

$$\mathbf{M} = \begin{pmatrix} N+n & N \\ N & N \end{pmatrix} \Longrightarrow \mathbf{M^{-1}} = \frac{1}{Nn}\begin{pmatrix} N & -N \\ -N & N+n \end{pmatrix}. \tag{31}$$

The inverse of the lower right element of $\mathbf{M^{-1}}$ is the effective number of observations or weight and is given by $n_{eff} = (Nn/(N+n))$, thus as $N \longrightarrow \infty$, $n_{eff} \longrightarrow n$. Therefore, no matter how precisely the sky spectra are known, the absolute sky is determined no better than the calibration.

7.2.2. β Uncertainties

The destriping effectively recalibrates the instrument eliminating the P_0EP_0 offset errors (and replacing them with the destriper errors). By fitting the sky with stripes, the offset errors in the emissivities of the *FIRAS* internal objects are eliminated. These are the P_0EP_0 offset errors. However the gain uncertainties remain.

The stripe determinations are not perfect and the uncertainty in their determination still leaves uncertainty in the sky maps. The HIGH data are destriped in three frequency segments (Section 6.3). These have varying numbers and types of stripes. We approximate all of the stripes with 10 stripes which are then good for all HIGH frequencies. There are two approximations in this proceedure. First, we concatinate the stripes for the three frequency segments weighted by the total weight of each frequency segment. The pixel weight is divided out of each pixel to give a fair comparison of the data. The matrix is then decomposed into eigen values and eigen vectors. Then we select the eigen vectors and values for the largest 10 eigen values. This set contains $\sim 99\%$ of the total power in the matrix. Finally the pixel weight is multiplied back to give the same form as the original matrix.

This is obviously a simplification, but since the original stripe kernels are empirically motivated, even the full set of stripes is a simplification. Further, by including the other segments' uncertainties we allow for variation in kernels that are used in the other segments, but do not have enough signal to be significant. Finally, the reduction to 10 stripes merely mathematically encodes what we already know: many combinations are well known, and the errors in them are unimportant. For example, the LHSS-RHSS is well determined, since virtually every LHSS coadd has an RHSS coadd to compare. However, since the HIGH uses the average, the LHSS+RHSS has the important uncertainty, although it is much less well determined.

For the combined LOWF there are 30 stripes. When these stripes are orthogonalized the largest few dominate the uncertainties. Therefore only the largest 10 need be used. This effectively summarizes our knowledge and uncertainty of the stripes that we applied. However, the stripes were largely motivated empirically, and the true form of the drifts remains hidden. For this reason we recommend including an uncertainty of $\sim .04\ C_\nu$ when making measurements on the scale of 10° and larger.

7.2.3. Gain Uncertainties

Another way of estimating the gain errors is to compare the results of different channels and scan modes (Section 7.7). The gain can be determined by weighting by the signal and the pixel weight of each channel and scan mode. This shows the gain variations have a random (in frequency) and correlated part. If the random part is identified as the P_gEP_g then it is larger by 1.7 than the P_gEP_g estimated from the calibration results. We recommend using this result because it is larger and it is derived from the final product.

7.3. The Bolometer Model Uncertainties

The bolometer model parameters that are fit in the calibration program have uncertainties which give rise to uncertainties in the sky spectra. These can be conveniently expressed collectively as the product of the Jacobian, J, and the covariance matrix, C, (not the C of Section 7.1) of the parameter uncertainties: $J \cdot C \cdot J^T$. A linear transformation to a different basis set is applied to the parameters to find the principal components of the covariance.

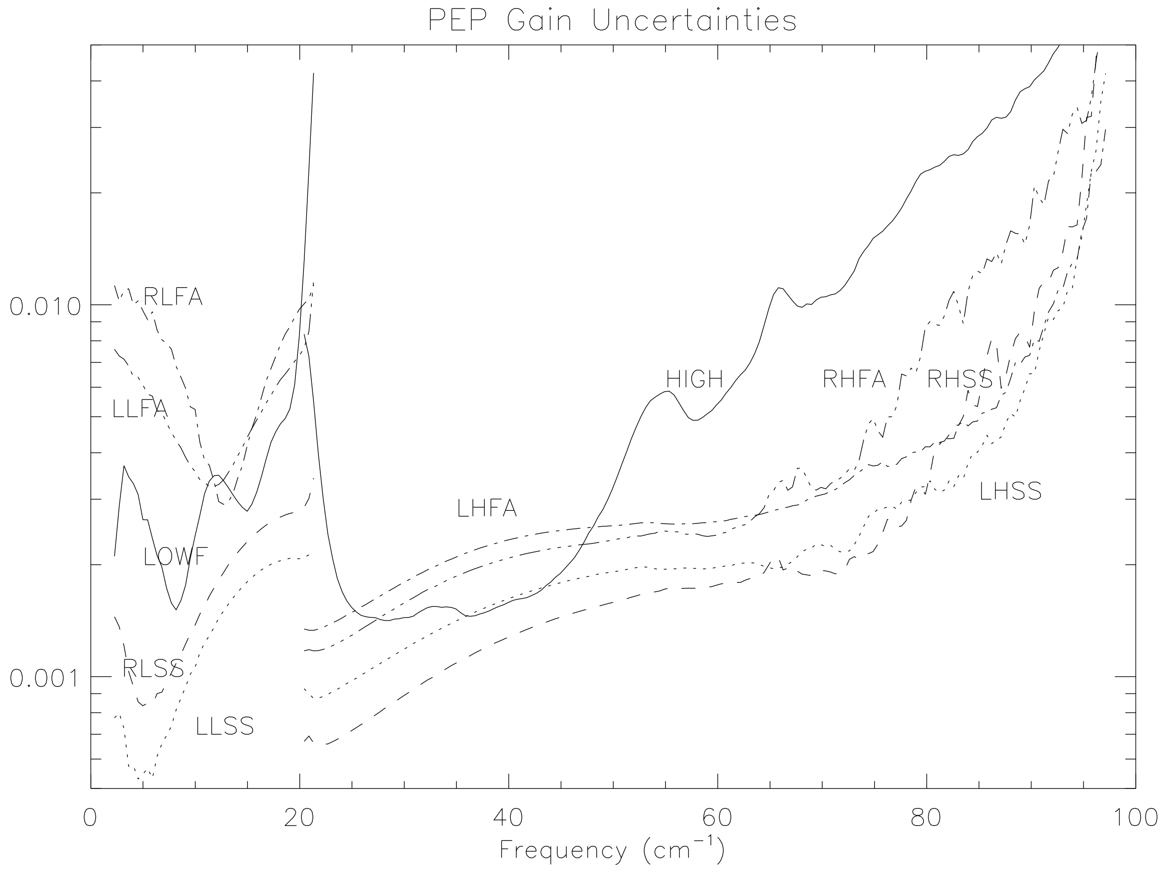

Fig. 7.3.— $P_g E P_g$ uncertainties — Estimated from the calibration models (LLSS, RLSS, LLFA, and RLFA) and the variation in the calibrated data. Gain uncertainty is dimensionless. The shape of the LOWF estimate is from the "C vector" and so should not be taken too seriously, but the scale is determined from the calibrated data variation.

Some errors can be well approximated as gain errors while others are better approximated as offset errors. To find these approximations the calibration parameters are perturbed and the difference in the calibrated sky is noted for 19 selected pixels. The pixels at high latitude, where the signal is small, are used to determine the "typical" offset while those at low latitude, where the signal is large, are used to determine the "typical" gain variation.

These uncertainty estimates are contained in the vectors *JCJ_OFFSET* and *JCJ_GAIN* fields of the uncertainty terms file. While all of the uncertainty terms are included for completeness, most of the uncertainty is included in a single uncertainty term.

7.3.1. JCJ Offset

Offset errors are corrected to first order by the destriper (and replaced by the destriper errors) so these are largely covered in the β matrix of the destriper. The β matrix is assumed to be independent of frequency, however, the *JCJ* offset errors are correlated. This issue will be addressed again below in the "FEF" errors. Detailed study has not shown the *JCJ* offset errors to be a problem either in theory or in fact for either the low frequency (where any *JCJ* is effectively overshadowed by the β) or the high frequency where the β errors are larger and the sky variation can mask the effects.

7.3.2. JCJ Gain

For many of the *JCJ* gain terms the largest effect is a broadband error. These are particularly significant for the high frequency data. The raw terms tend to be flat or slowly rising functions of frequency, sometimes with a steep increase at the high frequency end. This can be easily understood as the highest frequencies remain in the Wien portion of the spectrum even for the highest temperatures (~ 20 K), so a small change in these temperatures in the calibration can lead to a large change in the spectrum and hence in the gain. Over most of the spectrum however, there are at least some calibration data with the frequency in the Rayleigh-Jeans part of the spectrum. Over this range changes in the performance of the bolometer lead to changes in response over the whole spectrum, hence the slow functions of frequency. Even changes in the heat capacity or electrical capacitance effects lead to at most a linear dependence on frequency.

But *JCJ* gain uncertainties are a difficult issue. The formal solution covers the uncertainty in the model parameters, however it does not address the errors in the model $form$. The variance of the data is strongly dependent on the temperature, and higher temperatures

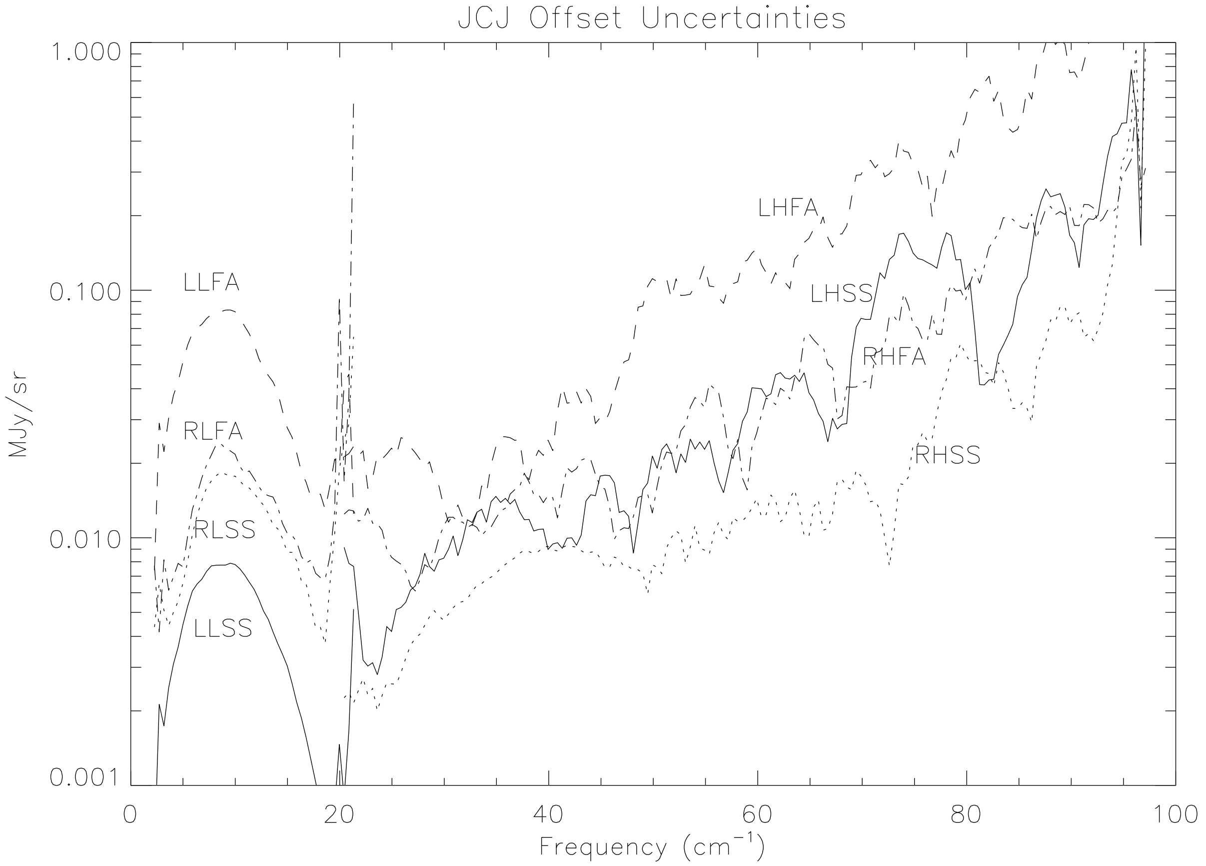

Fig. 7.4.— The *JCJ* offset uncertainties — Derived from the calibration models. The uncertainties reflect the amount and quality of the calibration data.

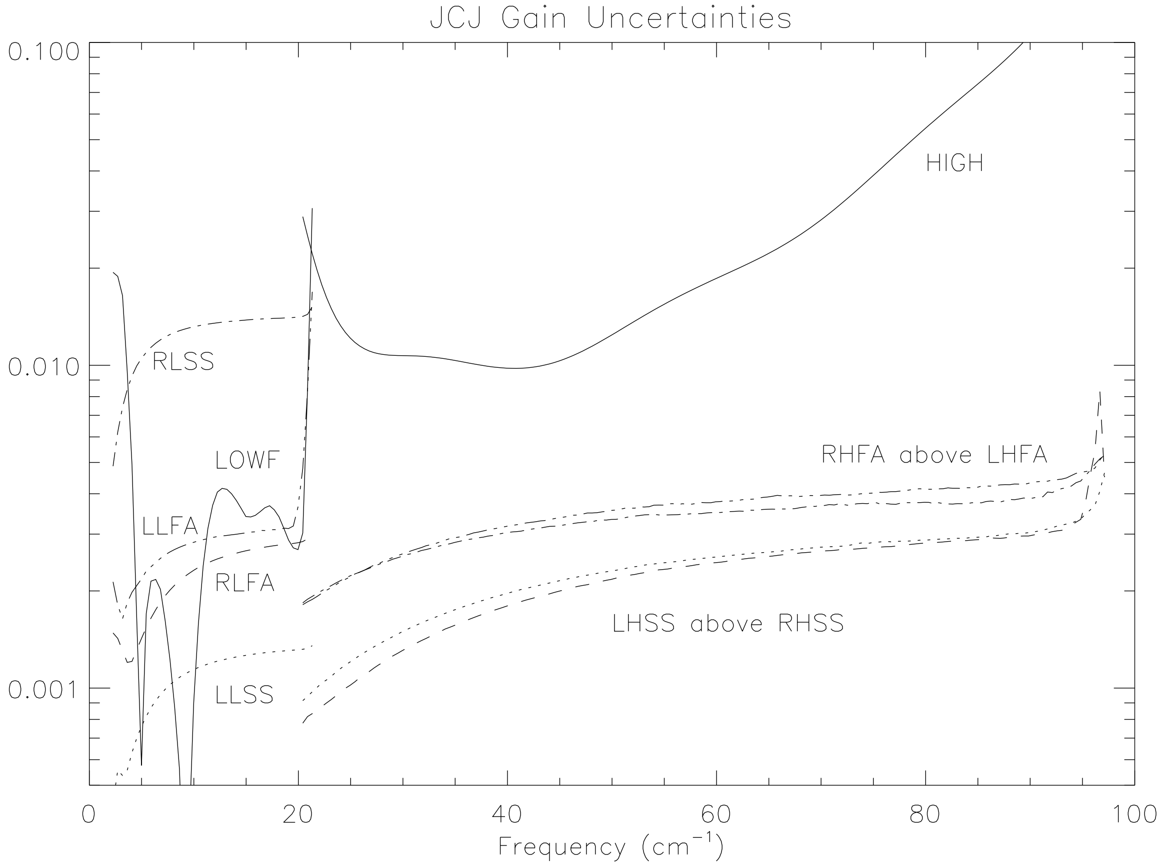

Fig. 7.5.— *JCJ* gain uncertainties — Derived from the calibration models except for LOWF and HIGH which are derived from the calibrated data differences. The dips in the LOWF are probably the result of a chance alignment of the gains.

lead to larger variance within the coadds. Even correcting for this, the χ^2 is larger for the high temperature data (Section 5.5). This is probably an indication of an inadequate bolometer model. Still, even a much more flexible bolometer model is likely to have similarly shaped Jacobians, because the major effects of the bolometer are in the responsivity and the time constant. Further the structure of the χ^2 shows only the features of the vibration lines and the transfer function. Since the radiative loading from the high temperature calibrations is much higher than any place on the sky, the calibration data stresses the model more than the sky data.

Barring the introduction of different parameters in the model we have only a limited means to estimate the true uncertainties. One method is to compare different channels. The random variations were identified as P_gEP_g errors. Identifying the correlated variations in gain as *JCJ* errors leads to the conclusion that the calibration *JCJ* determination underestimates the *JCJ* gain error by a factor of 5. This suggests errors of 1%, however, the basic model is the same in all cases. A missing parameter "X" could well have a similar effect for all channels. If such a factor exists, it probably depends either on the radiative loading on the dectector or on the time constant of the detector. The implied power is different by a factor of ~ 4 between the left and right and the output frequency is in the ratio 3:2 for fast and slow. So, although "X" may push all of the channels in the same direction, it should push left or right ~ 4 times further or push fast or slow 1.5 times further. In either case, the correlated gain errors are the right scale. The results from FISH typically estimate the high frequency gain uncertainties a factor of 2-4 larger than the low frequency uncertainties, and, in fact, they appear a factor of ~ 4 larger. We recommend the combined *JCJ* gain uncertainties plotted in Figure 7.5 or 0.4% for LOWF.

The differences between LLLF and RLLF are too noisy to use to estimate the gain uncertainty. The HRES P_0EP_0 uncertainty is estimated from the FISH estimates by multiplying by 1.7, the ratio observed in LOWF. The HRES *JCJ* estimate is just two times the LOWF estimate. The factor of two comes from the ratio of the amount of calibration data in HRES to that of LOWF. The correlated gain errors are treated next.

7.4. The Temperature Noise

Errors in the temperature of the various components of the *FIRAS* while observing the sky get wrongly attributed to the sky. Each error δT_k of course is multiplied by the emissivity of object e_k and the Planck function, $P_k = \partial \mathrm{Planck}/\partial T|_{T_k}$. The errors and emissivities can be summed into a convenient matrix $U_{kk'}$. So the variance can be expressed as $\sum_{kk'} P_{k\nu} U_{kk'pp'} P_{k'\nu'}$. We assume that the off diagonal terms (that is the correlations between thermometers) are negligible, but it is not required for the formalism. The emissivities show that U is dominated by the ICAL, $e_I \sim .98$ (the XCAL is not in the horn for sky observations), except during the "hot horn" season which is treated separately. Since all of the critical elements are kept at 2.7 K for most of the flight the subscript of the P_k is dropped. If we make the assumption (not entirely warranted) that the ICAL temperature readings are independent then the pp' part of the U becomes diagonal and can be conveniently summarized as $1/N_p$ where N_p is the number of observations of the pixel. Then the PUP problem is reduced to estimating a single constant: the uncertainty of reading the ICAL temperature in an IFG. There are three separate ways of estimating this uncertainty.

7.4.1. Resistor Noise

To calibrate the Germanium Resistance Thermometers (GRTs), *FIRAS* uses a set of 4 fixed resistors in the circuit. These were measured along with the GRTs every 32 seconds during the 10 month mission. These $\sim$ 3 million measurements can also be used to estimate the readout noise of the electronics *in situ* to high precision. The rms estimate, 16.6 bits is good to 1% for all of the resistors. Multiplying this by an estimate of 130 μK per bit gives an rms uncertainty of 2158 μK per readout.

Since the critical ICAL temperatures are derived from two thermometers on two sides as the interpolation between two times, the temperatures are derived from eight separate readouts. This implies an rms uncertainty of 880 μK on each IFG and a similar uncertainty on each sky measurement.

This estimate ignores possible correlations and effects that might make the GRT readings less stable than the resistor measurements, so this can be treated as a lower limit to the PUP uncertainty.

7.4.2. Calibration Noise

In the calibration program (FISH) we explicitly solve for the temperatures of the ICAL, XCAL etc. This list of corrections can be used as an estimator of the rms errors of the temperature readout and hence an estimator of the PUP uncertainty.

The rms weighted noise estimated in this way is 837 μK per IFG. However there are two corrections which must be applied to get an estimate of the PUP uncertainty. First it must be multiplied by $\sqrt{2}$ because FISH is only sensitive to the difference in temperatures in the cold data set (all controllables between 2.6 K and 2.8 K), which is divided between the ICAL and XCAL corrections. Second in the calibration data the IFGs are taken "back to back" so there are correlations introduced by the interpolation to get the temperatures. On average these inflate the final noise by 1.09. So after correcting for these effects the estimated PUP uncertainty is 1086 μK .

At first blush it might seem that this is a perfect estimator of the PUP uncertainty, however there are stripe errors (which we later find and remove in the destriper) and these are included in the estimate. Also there are only of order 300 cold coadds from which to make the estimate. The first item means that this estimator may be biased high and the second suggests uncertainties in the PUP estimate of $\sim 3\%$.

7.4.3. Resulting Noise

The final answer is always a good place to look for noise. In this case we can use several techniques to extract the CMB temperature (Fixsen 1997a) and that temperature can be compared to the known temperature. In the comparison to the *DMR* temperatures we note there is "extra noise" in the *FIRAS* data at ~ 45 μK per *DMR* beam not accounted for by other uncertainties. On average there are 767 IFGs/beam. A *DMR* beam is 16.4 pixels. There are 126382 LLSS IFGs and 157207 LLFA IFGs in the sky data. The RL data are not counted, as their PUP error is fully correlated with the corresponding LL PUP error. This gives an estimate of 1246 μK per IFG.

Since this is an estimate of PUP from the final data it leaves little chance that there is something else "hiding" in the data. However, this discrepancy may be due, at least in part, to other errors. In particular the stripe errors may make a significant contribution to this uncertainty. Still this provides a firm upper limit to the PUP uncertainty.

7.4.4. Hot Horns

For 68 days during the mission the horns (both the sky horn and the reference horn) were held at an elevated temperature, first 6K and then 4K. Although the horn emissivities are much smaller than the ICAL the higher temperatures more than overcome the emissivity and the the net PUP is larger than during normal times. However the PUP in this case is $\nu dP/dT$ at 6K rather than dP/dT at 2.7K, because the horn $e \propto \nu$.

7.4.5. *Recommended PUP*

The three quite different methods of estimating the PUP uncertainty vary by only 30%. Further the order is what is expected given the possibility of contamination by other uncertainties. We recommend the use of 1 mK per IFG or 150 μK per observation as a good estimate of the PUP uncertainty, with an additional 8μK $\nu dP/dT$ at 6K during the hot horn time.

7.5. The Temperature Uncertainty

Thermometry errors during calibration induce systematic errors in the calibration model parameters. These are collectively called the PTP error. Since there are a large number of measurements in the calibration data ($\sim$ 20000), small random errors largely cancel out. Furthermore, systematic errors in the thermometers of the internal calibrator, the horns, and other internal components of the instrument will be repeated while taking sky data. Hence, the "errors" induced in the model will later be removed by the same "errors" when observing the sky. This is not so for the external calibrator, for which any systematic errors are reproduced in the final sky spectra. Thus, the dominant contribution to the PTP uncertainty is the systematic error in the external calibrator. The form of the PTP error is

$$\frac{\partial P}{\partial T} \cdot T \cdot \frac{\partial P}{\partial T} \tag{32}$$

where the outer factor is the partial derivative of the Planck function of the CMBR with respect to temperature, and T is the variance of the thermometer uncertainty for XCAL. The estimate of the absolute thermometry uncertainty is based primarily on the disagreement between the photometric and thermometric models, which differed by $\sim$4.5 mK at 2.7 K. We have split the difference and adopted 2 mK as the one σ uncertainty in the XCAL thermometer. This uncertainty estimate is contained in the PTP_TEMP and PTP_SPEC fields of the uncertainty term file. These were checked by a third method which uses the Dipole of the CMB. This uses the same frequency scale; however, if we use the *DMR* to recalibrate the amplitude, we have an independent check, which gives a correction of -7 ± 9mK to the XCAL thermometry.

These three methods give answers that are within three sigma. Apparently there are additional systematic errors in one or more of these methods. We recommend a -2.25 mK temperature adjustment to the XCAL diode temperature scale. This has been applied to the calibrated data. We recommend a 2 mK estimate of the PTP uncertainty. This allows covering both the color temperature calibration and the thermometer calibration. While this is not a true statistical uncertainty it is a useful summary of the uncertainty in the result. The PTP error is the dominant error for the absolute temperature of the CMBR and is important for comparisons of *FIRAS* measurements to other experiments.

7.6. Full Uncertainties

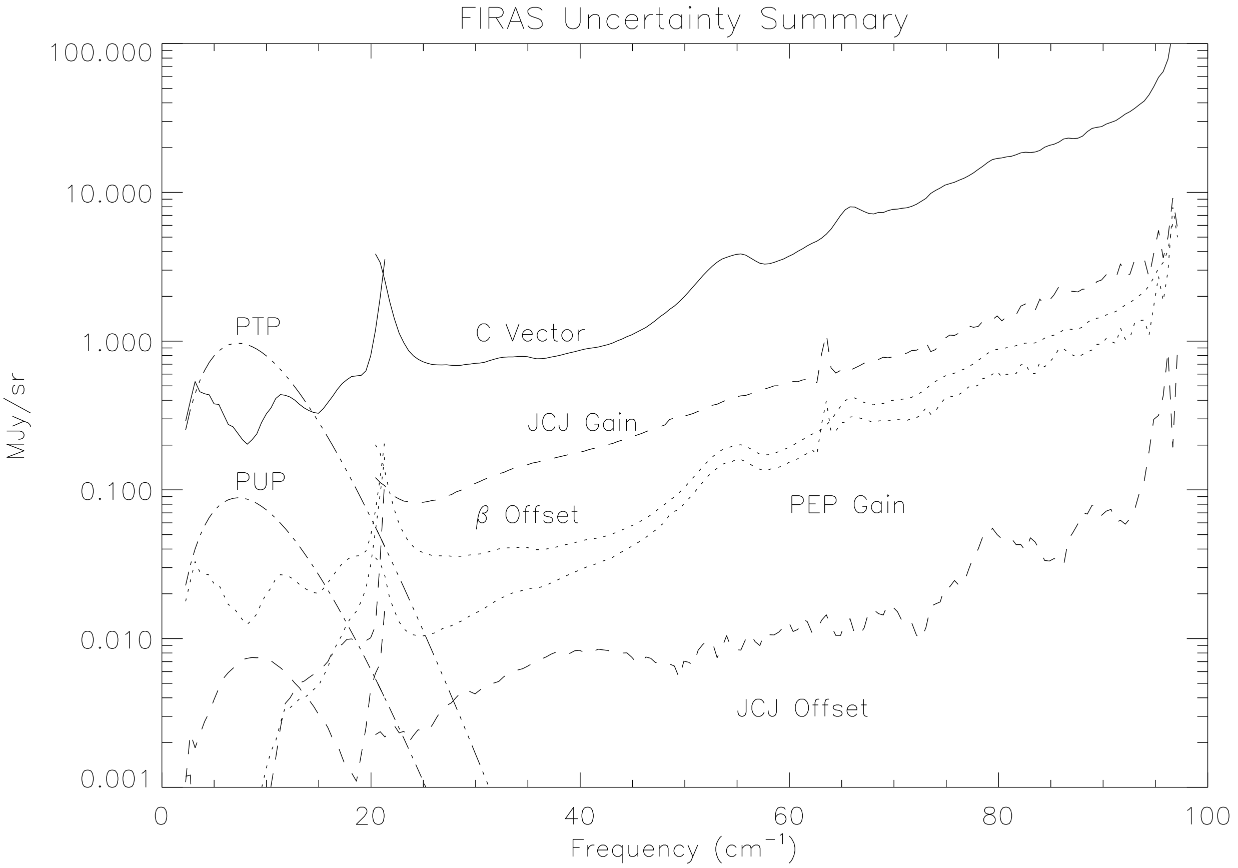

Fig. 7.6.— *FIRAS* Uncertainty Summary — Plots of calibration model solution uncertainty estimates from the C vector, the P_0EP_0, the β, the JCJ, the PUP, and the PTP uncertainties. The uncertainties for JCJ gain and P_gEP_g have been multiplied by the average sky.

The calibration model uncertainties are shown in Figure 7.6. The offset uncertainties have contributions from all five error terms, while the gain uncertainties only have contributions from the P_gEP_g and JCJ uncertainties. The gain uncertainties are multiplied by the average sky in order to make a meaningful comparison. Examination of the offset uncertainties shows that for analyses using less then ∼200 observations, the *C-Vector* uncertainties dominate, but for larger data sets other uncertainties must be taken into consideration. Examination of the gain uncertainties shows that the *JCJ* gain error dominates the P_gEP_g uncertainties. The effect of these uncertainties on particular analyses is discussed in Section 7.10.

The largest uncertainty (at least over part of the range) is the *PTP* so when it applies, that is, when the question is what is the absolute temperature, it is the dominant term. Next is the *C-Vector*, but this is for a single average pixel. By averaging or fitting over either pixels or frequencies it can be suppressed. When averaging or fitting over pixels the *PUP* is also suppressed, so it is only important when finding temperature variations and then it is smaller although on the the same order as the *C-Vector*. The *JCJ* gain uncertainty is important at higher frequencies particularly if fitting or averaging over frequencies. The β offset takes the place of the P_0EP_0 which dominates when averaging more than about 400 "observations". The β shown here is an "average" over the sky; it can be larger or smaller depending on the sky sample. The P_gEP_g uncertainty is only important when looking at differential measurements over frequency, e.g. finding lines. We have not found the *JCJ* offset to be a significant uncertainty but we include it here for completeness.

7.7. χ^2 Distributions of Combined Skymaps

We have compared the actual χ^2 distributions to expected distributions for HIGH, LOWF, and HRES. In the HIGH and LOWF combinations the Galaxy stands out as a region of high χ^2. This is due to the fast and slow data sets being taken at different times and hence the registration difference is of the order of a pixel (2.5°). This has only a small effect where there is only a slowly varying signal, but near the Galactic plane it is significant. The errors show the largest χ^2 is not on the Galactic plane but approximately 3° to either side where the beam is just starting to pick up the Galactic plane.

The HIGH skymap displays an excess in χ^2 due to position variation at Galactic latitudes less than 10° (Figure 7.7). The effect is smaller for the LOWF because the high frequency band is more sensitive to the Galactic signal. As with LOWF, sets of IFGs for HIGH were obtained at different times.

The combined χ^2 *vs.* frequency for HIGH averaged for all pixels greater than 8° from the Galactic plane shows the χ^2 to be anomalously large at the seven highest frequencies. The actual and expected χ^2 distributions agree well, however, if one assumes an 8% random

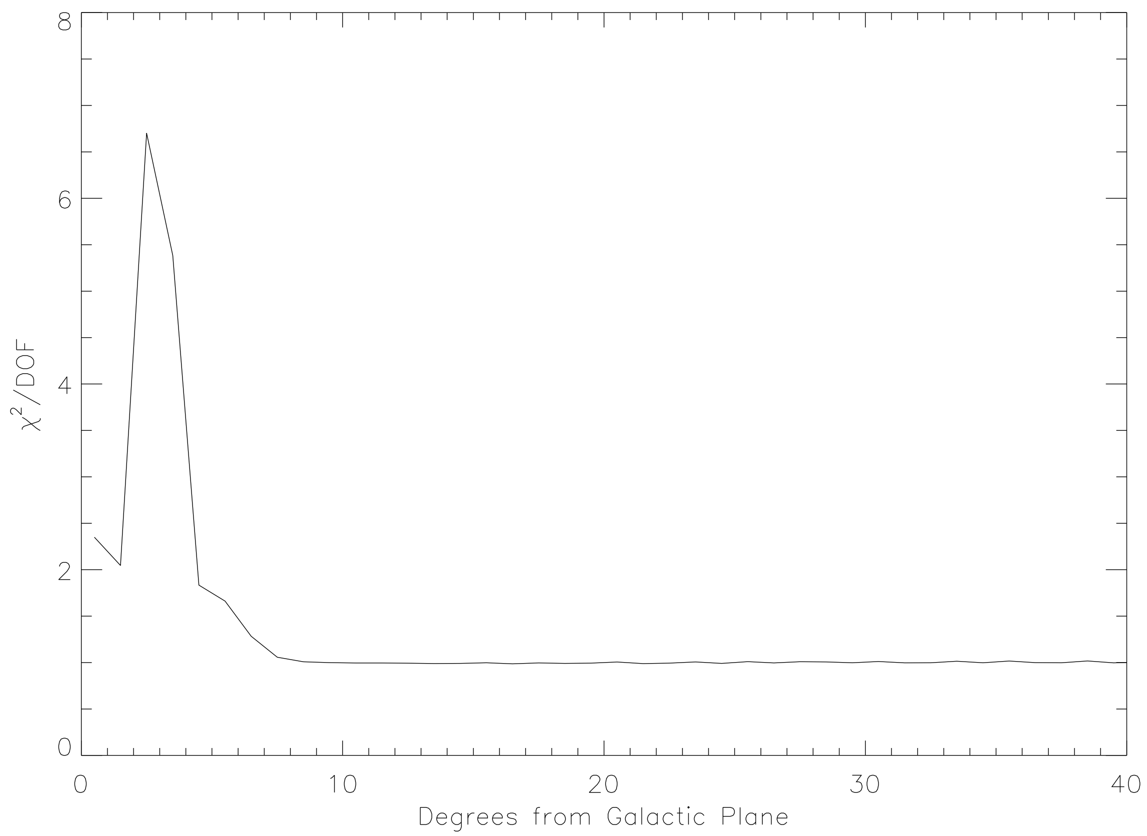

Fig. 7.7.— χ^2/DOF as a function of Galactic latitude — The gradient induces variation in the data larger than the intrinsic variance near the Galactic plane.

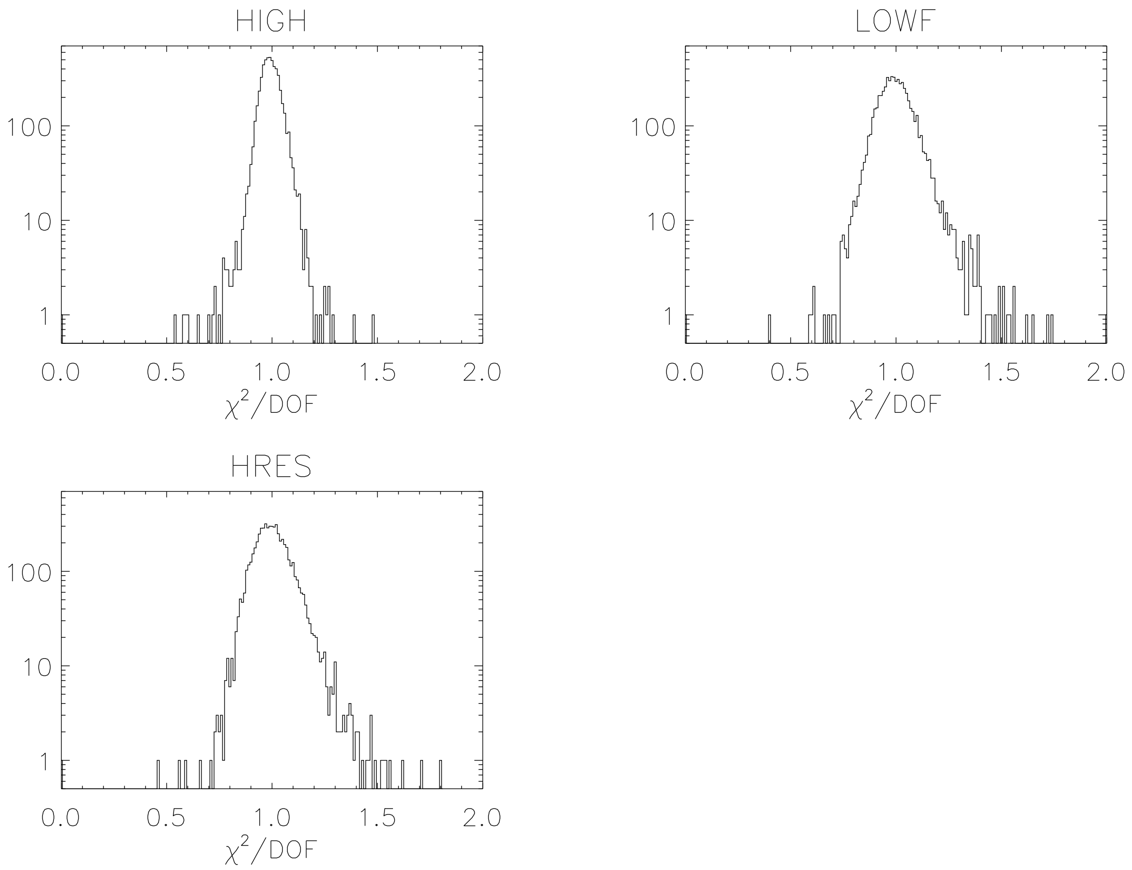

Fig. 7.8.— Histograms of the χ^2/DOF of the HIGH, LOWF, and HRES combinations — The DOF for these combinations are 170, 43, and 182, respectively. The Galactic center is omitted from these plots.

error in the pixel weight (Figure 7.8).

The combined χ^2 averaged for LOWF for all pixels greater than 5°away from the Galactic plane has a broader distribution than expected. The actual and expected χ^2 distributions agree well, if one assumes a 15% random error in the pixel weight (Figure 7.8).

The HRES map shows no evidence of the Galaxy. This is expected, since the input LLLF and RLLF maps have minimal registration difference. The HRES map does display stripes of anomalously low and high χ^2 . The relation of these stripes to mission intervals when the sky and reference horns were raised to higher temperatures suggests that errors in the calibration model for LLLF and/or RLLF may not have been entirely removed by destriping. The other maps show high χ^2 associated with the Galaxy, but are otherwise free of obvious systematic effects.

7.8. "*FIRAS* Extra Factor" Errors

A way to investigate possible errors is to change the processing in "modest and reasonable" ways and to note the variation in the output. Although not particularly quantitative, these variations indicate the type and scale of errors introduced by the restrictions of the models used in processing the *FIRAS* data. These *FIRAS* Extra Factors (FEF) are not rigorous errors in the statistical sense, but are included to point out possible pitfalls for users of the *FIRAS* data.

7.8.1. The Destriper Model

The final results from the *FIRAS* data depend on several choices concerning the destriper model. The choices made in the data processing were the "best" choices in our estimation. Other choices are possible, so we have provided the change in offset for several different parameterizations (Figures 7.9 to 7.14). These changes can be used to estimate an uncertainty in a final result. This is not straight forward, since a one σ result is not defined for a change in the model. Rather than a true estimation of errors, these changes should be used to allay fears (in many cases) or to emphasize caution. In all cases, we have estimated an offset by subtracting the nominal result from a variant averaged over $|b| > 60°$, which measures the offset over a cosmologically interesting region of the sky.

The nominal destriper model has from 10 to 30 terms (Section 6.2). The project data sets contain five sets of comparisons of destriper models: three frequency band comparisons, a low frequency low resolution comparison, and a low frequency high resolution comparison.

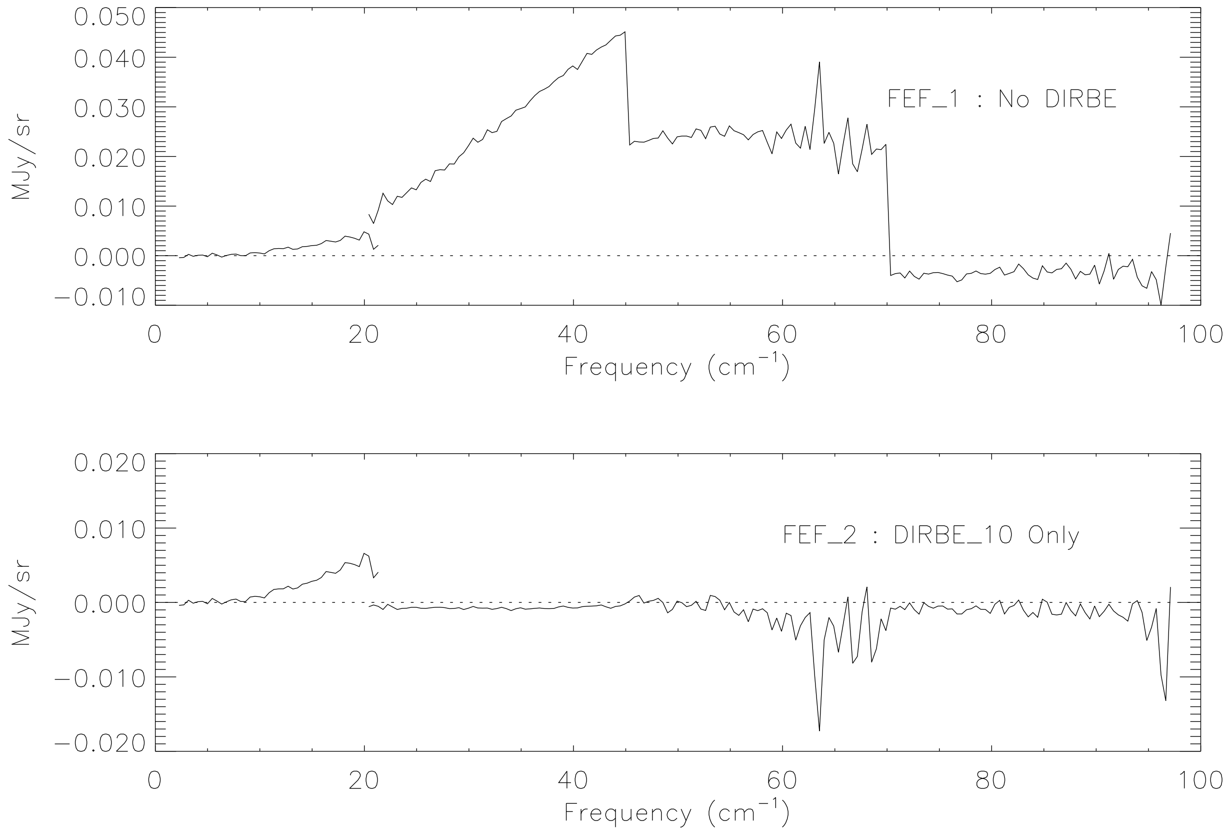

Fig. 7.9.— Galactic gradient model difference spectra — Average variations in the sky spectra for changes in the number or form of the gradient models.

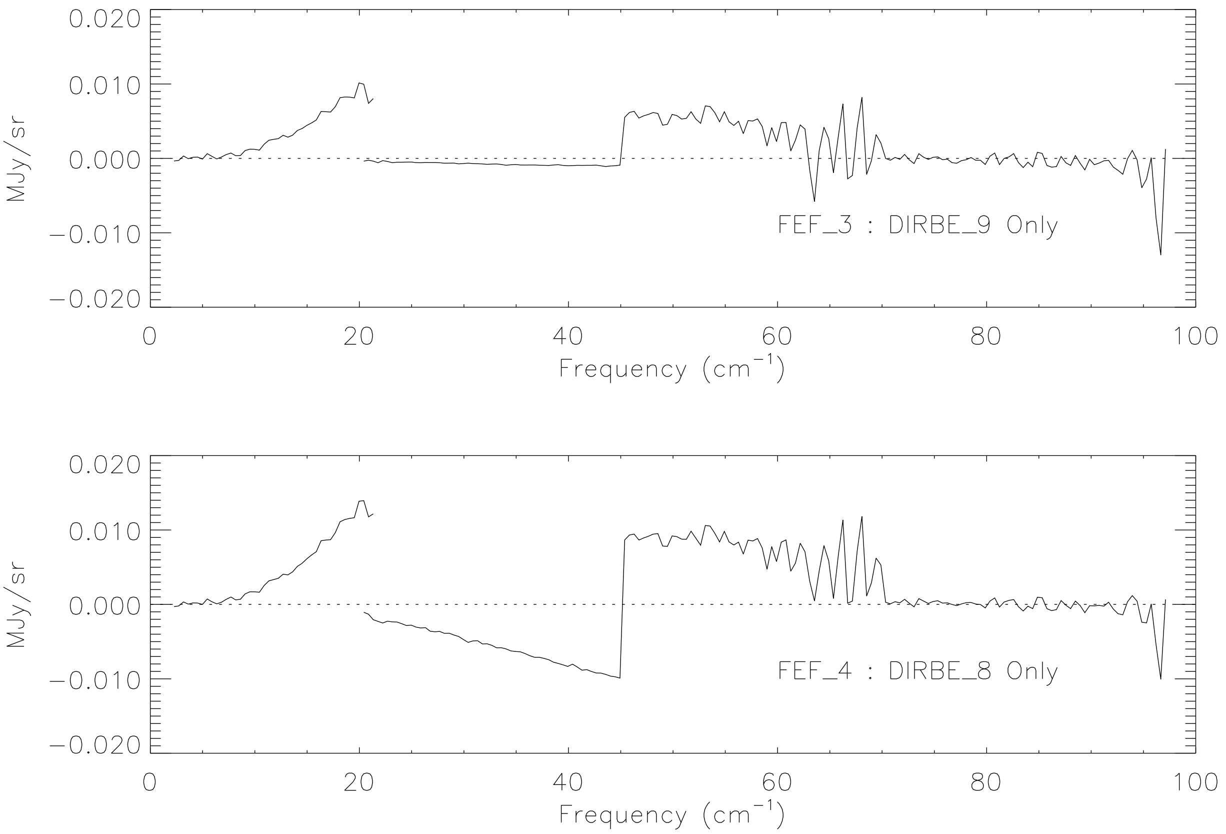

Fig. 7.10.— Galactic gradient model difference spectra (continued) — Q.v. Figure 7.9

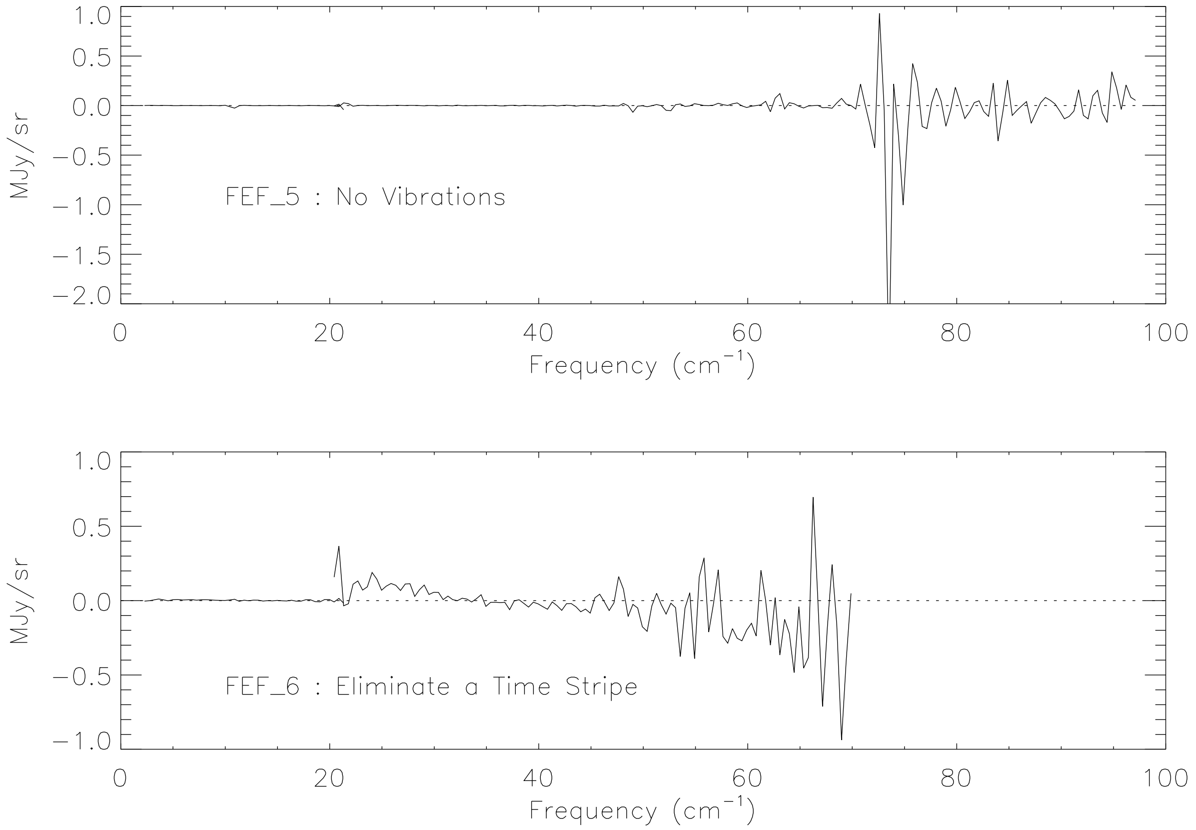

Fig. 7.11.— Vibration and time stripe difference spectra — Average variations in the sky spectra if no vibration stripe was included or if one less time stripe was included. Even though the effect is small in the LOWF, it is still significant.

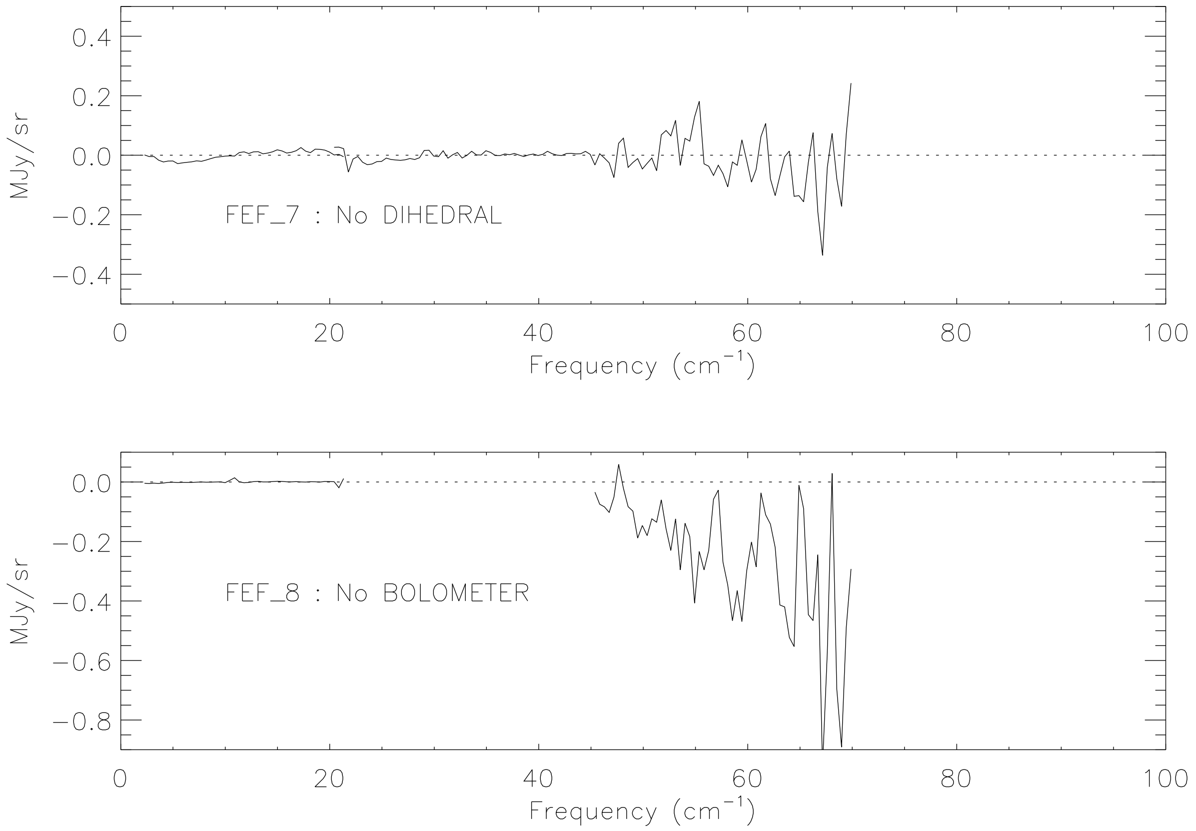

Fig. 7.12.— Dihedral and bolometer stripe difference spectra — Average variations in the sky spectra if the dihedral and bolometer stripes were not included. The effects in LOWF are significant, if small.

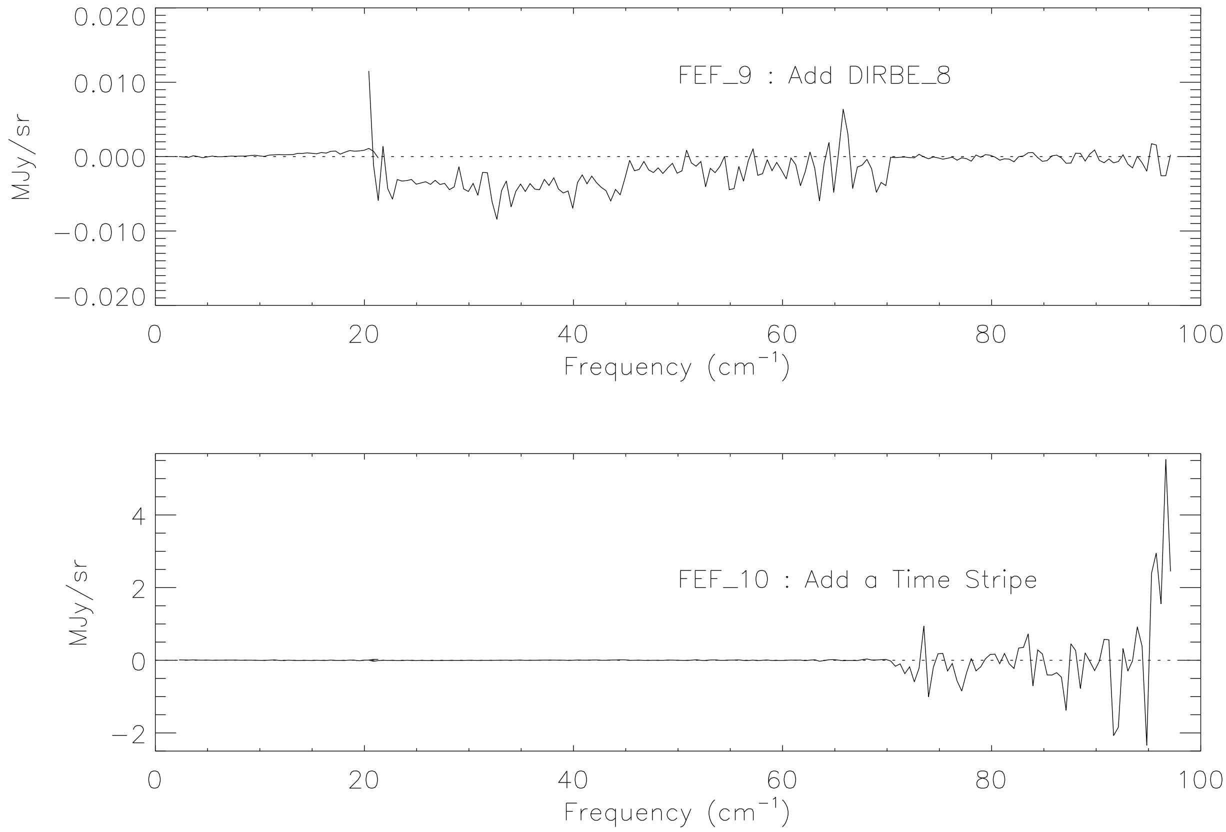

Fig. 7.13.— Additional galactic gradient and time model difference spectra — Average variations in the sky spectra when an additional DIRBE based stripe or an additional time stripe are used. The magnitude of the time stripe difference appears large at high frequencies, but the uncertainties are large there as well. The stripe is not significant by the F test.

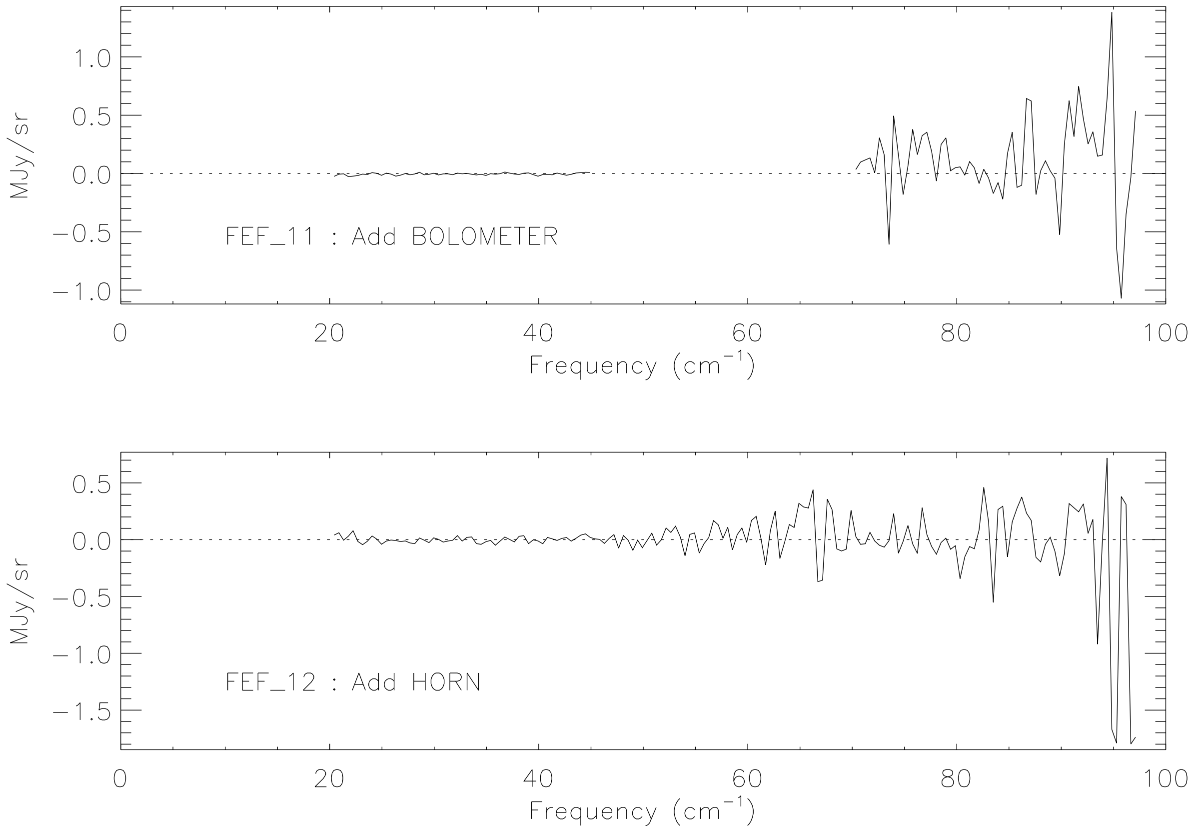

Fig. 7.14.— Additional bolometer and 'hot horn' difference spectra — Average variations in the sky spectra when a bolometer stripe, a 4K horn, or 6K horn stripe, for data where such stripes were not used. The additions fail the F test.

Table 7.1: *FEF* Comparisons

Comparison	Description
FEF_1	No *DIRBE*
FEF_2	*DIRBE* 10 Only
FEF_3	*DIRBE* 9 Only
FEF_4	*DIRBE* 8 Only
FEF_5	No Vibration
FEF_6	Eliminating a Time Stripe
FEF_7	No Dihedral
FEF_8	No Bolometer
FEF_9	Three *DIRBE* Channels (Adding Band 8)
FEF_10	Adding a Time Stripe
FEF_11	Add Bolometer
FEF_12	Add Horn

Each of these data sets contains comparisons of destriper models, which are shown in Table 7.1. The comparisons for LOWF and HIGH are shown in Figures 7.9 through 7.14. The HRES comparisons are available, but are not plotted here.

Together, these comparisons show that final results are quite robust under variations of the destriper model. The LOWF variations are typically on the scale of a few kJy/sr ($\sim 10^{-5}$ of the CMBR peak). These are $< 0.1\ \sigma$ for a large fraction of the data, so we can be confident that any reasonable destriper will give a comparable answer. Since the same data are going into both destripers the errors could in principle be zero. A point at 10.6 cm^{-1} (318 GHz) often sticks out in the LLSS and RLSS data; this is the vibration discussed in Sections 6.3.7 and 7.9.1.

More problems are visible in the high frequency data. The vibration stands out at 73.5 cm^{-1} (2203 GHz) and 49 cm^{-1} (1469 GHz) in the slow and fast data respectively. The high frequency ($>$ 70 cm^{-1}or 2100 GHz) fall off in the performance of the interferometer is evident in the larger errors there. Still the high frequency data are robust under the various choices of destriper, with variations $< 0.1\ \sigma$.

Since the gain variation is an issue for the high frequency channels (Fixsen *et al.*, 1994b) we have included the gain under changes in the destriper model. Not surprisingly the gain changes show little variation since the destriper does not adjust the gain.

7.8.2. Channel Consistency

One of the strengths of the *FIRAS* is that it has multiple channels which cover the same parts of the spectrum. These allow cross checks on the results at many points. If we look at the differences between the channels we see they generally agree to within a few σ. These differences are statistically significant, but we have no clear explanation of their causes. This test is not as powerful as one might hope since the LH detector has much more noise than the RH detector and the RL detector has more noise than the LL detector.

Comparing the gain of the various channels also indicates the possibility of problems. The gain ratio test is only valid for $\nu >\sim 25$ cm^{-1} (750 GHz) since absolute spectra were compared rather than differential spectra (Section 7.3).

To estimate the variations in gain, each channel combination was used to construct a separate map. Each of these maps was then averaged with a weight that included a factor for the noise of each of the maps and a factor for the amplitude of the signal. The result is a spectrum for each of the channels, and a comparison spectrum for the full map. The ratios of these should ideally be one, however, noise and gain errors will induce variations. The variations are fit with a polynomial (10th order for HIGH and 8th order for LOWF). The polynomials are identified as the systematic variations (*JCJ*), while the residuals are identified as the random (in frequency) variations (P_gEP_g errors). These are larger than expected by the gain uncertainties propagated from the FISH results, particularly in the case of *JCJ*. So, the results were combined to give an estimate for the P_gEP_g and *JCJ* uncertainties (Sections 7.2.3, 7.3.2) in the final HIGH and LOWF data sets. Comparison noise dominates in the HRES case.

We have compared the *FIRAS* HIGH data to the *DIRBE* bands 8, 9, and 10 data. This is a robust check since the *DIRBE* instrument is quite different, has a different calibration standard, and looks at the same region of the sky at different times from the *FIRAS* observations. The results for band 10 ($\nu = 42 \pm 8$ cm^{-1}) and band 9 ($\nu = 71 \pm 10$ cm^{-1}) show that to within $\sim$ 5% *DIRBE* and *FIRAS* HIGH data agree. While the absolute *DIRBE* gain calibration is only known to 10%, the agreement with *FIRAS* suggests that both are correct to ∼5%. At band 8 ($\nu = 100 \pm 20$ cm^{-1}) the agreement is good to 20%. Noise on the *FIRAS* data and the necessity of a color correction limit the accuracy of the band 8 comparison.

The comparison with *DMR* (Fixsen *et al.* 1997a) gives less than 0.5% difference in the amplitude of the dipole. This suggests that the gain error is <0.5% for the low frequencies, in agreement with our estimate of 0.4%..

7.9. Unmodeled Systematic Errors

In addition to the modeled uncertainties other discrepancies have appeared in the data analysis. These problems are difficult to quantify and their consideration can be awkward but in some applications they dominate the error picture and so they must be considered.

7.9.1. Vibration

There are major MTM resonances at 57.57 Hz and 8.32 Hz, and there are effects on the spectra at harmonics and sum and difference frequencies associated with these resonances. A vibration correction term in the calibration model corrects most effects. However, at the vibration frequencies themselves, the data cannot be entirely corrected. Destriping has greatly suppressed this problem, but one should be wary of any result that depends on the vibration frequencies (Section 6.3.7). The data frequencies translate to optical frequencies 73.5 and 10.6 cm^{-1} (2203 and 318 GHz) for slow scans or 49.0 and 7.1 cm^{-1} (1469 and 212 GHz) for fast scans.

For the high fast data the second harmonic of the vibration (115.14 Hz) occurs at an optical frequency of 98.0 cm^{-1}. In addition, there are sidebands of this harmonic that can contaminate the data above 90 cm^{-1}. These sidebands have not been modeled in the calibration, but have been suppressed by destriping. The second harmonic should be small relative to the main vibration but the signals at high frequencies are also small. In the high frequency slow data the second harmonic is pushed to 147 cm^{-1} and even the sidebands will be strongly attenuated by 120 cm^{-1} so the slow data are not contaminated in this way.

7.9.2. Phase Corrections

There are some significant phase shifts in some of the data. We provide undestriped complex spectra as they are derived from the Fourier transformation. Ideally, the imaginary parts of these spectra would be zero except in regions of the spectra where there are high resolution spectral features, *e.g.*, lines or bandpass filter edges. The line-like features in the imaginary parts are important in the determination of flux and frequency of real astrophysical line emission, but the wideband features in the imaginary parts are presumed to indicate instrument or data analysis faults.

The principal sources of phase variation in the complex spectra are resonant vibrations in the MTM and a small phase shift of unknown origin that is largely linear in frequency (*i.e.*

equivalent to a small displacement of the interferogram).

In addition to the MTM, an unknown source induced both short term and long term interferogram phase variations. It is important for high signal (*i.e.* far from null) data, which occurs primarily in calibration and in the high frequency channels in the Galactic plane. To account for this, we have introduced a phase corrector of the form $e^{i\nu\psi}$, where ψ is typically $< 10^{-4}$ cm but is statistically significant. Because the phase is time variable we fit a ψ to each calibration spectrum. It is not possible to follow an identical phase correction procedure for "cold" sky data, most of which are near null, so an analogous correction is applied wherein the phase is determined by minimizing the imaginary component of the sky spectrum. When applied to the calibration data as well, the results of this algorithm match the results of the more sophisticated fit. Phase errors are most important for the imaginary parts of the spectrum because $e^{i\nu\psi} \sim (1 - (\nu\psi)^2/2 + i\nu\psi)$ for small errors. The effect on the imaginary part is linear in the error, while the effect on the real part is second order as most of the information is in the real part of the spectrum.

7.9.3. Pointing Errors

The pointing solution is from the *DIRBE* data which have uncertainties smaller than 0.1° (Wright 1990). Since the *FIRAS* beam is by definition aligned with the spin axis of the *COBE* spacecraft, the *FIRAS* absolute and relative pointing is good to this same level. This is much better than is required for the 7° beam.

Each IFG is assigned to a pixel according to the line of sight at the center of the observing time interval, but the line of sight moves during that time. In addition, each individual IFG in a coadd group will have sampled its own distinct swath of the sky about that pixel. The individual swaths will largely overlap, since the pixel is oversampled, but the overlap will not be exact. In regions of the sky with large gradients in brightness across the field of view, especially the Galactic plane, we can expect the mean measured spectrum for a pixel to depend upon the details of the pixel sampling (Section 7.6).

The data preparation averages all the observations within a given pixel that are sufficiently consistent with one another and computes a mean line of sight for these consistent interferograms. The process of consistency determination is nonlinear, since it is intended to eliminate seriously defective IFGs. It is therefore difficult to make a quantitative estimate of these errors from the dispersion of the position coordinates alone.

A reasonably reliable estimate can be made from the dispersion of the spectra of individual coadds from the pixel mean. For most pixels, this dispersion is dominated by the random noise described by the "C Vector". Near the Galactic plane, the pointing dispersion dominates (Figure 7.7).

7.9.4. The Beam Profile

The sky horn antenna of the *FIRAS* is imperfect, especially at the higher frequencies. Observations of the Moon show both radial and azimuthal deviations from the nominal 7° circular top hat beam profile (Figure 7.15).

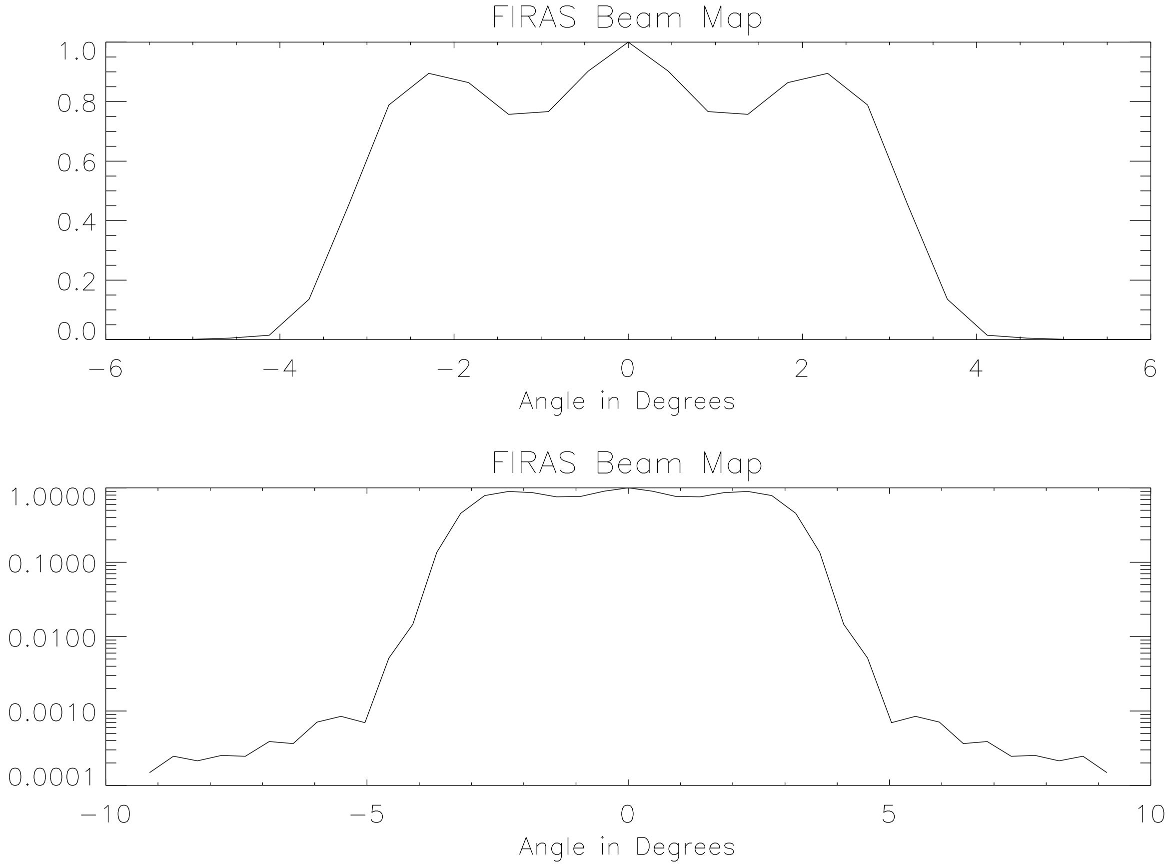

Fig. 7.15.— *FIRAS* beam map — Determined by moon observations. Both linear and logarithmic plots are displayed to show the structure in the beam and the sidelobes.

The measured profile is uniform with respect to radial angle for lunar aspect angles less than 2.8° . About 73% of a 25 cm^{-1} point source flux is contained within a circle of 5.6° diameter. An additional 24% of the flux is contained within an annulus of 5.6° to 7° , with about 3% leaking out of the nominal beam. Within the central 5.6° circle, we observe azimuthal variations of 35% at 25 cm^{-1} and 65% at 70 cm^{-1}.

Although these variations appear to be large, it is important to note that each sky pixel contains data taken from on average 100 IFGs, and each of those is the average of 16 strokes of the interferometer taken while the spacecraft spins about half a revolution and moves its line of sight. Therefore, the effective beam profile is the average of many profiles taken in a large number of possible orientations and positions, and the residual errors are expected to be very small. It is sufficient to note that the beam profile is reflected in the pointing variance.

7.10. Propagation of Errors

The data can be thought of as a matrix with each row as a map at a particular frequency, and each column as a spectrum of a particular pixel. Thus the data can be labeled $S_{p\nu}$. A full covariance matrix would be of the form $V_{pp'\nu\nu'}$ which has $(210 \cdot 6144)^2 \sim 10^{12}$ components. This is too large to store or use easily, so we store five parts which must be added to get the covariance matrix:

$$V \;=\; D\,(or\ C) \;+\; P_0EP_0 \;+\; JCJ \;+\; PUP \;+\; PTP \tag{33}$$

Since the covariances of each uncertainty term add linearly, the uncertainty vectors add in quadrature. Normally, only one or two of the dominant uncertainties need be considered when analyzing the data, but which terms dominate depends on how the data are used. In general, averaging the spectra over different parameters will reduce one or more of the uncertainties leaving the others dominant.

The effect of averaging or fitting spectra across pixels reduces the D and PUP uncertainties. The effect of averaging or fitting spectra across frequencies reduces the D and P_0EP_0 uncertainties. The PTP and JCJ uncertainties remain unchanged no matter how the data are averaged.

To explicitly state the full propagation of uncertainties, suppose a solution of interest, A^a, can be obtained from the *FIRAS* data $S^{p\nu}$ by some linear operator H:

$$A^a = H^a_{p\nu}S^{p\nu} \tag{34}$$

where we have adopted the Einstein summation convention. The index p runs over the list of pixels and ν runs over the frequencies. Now the uncertainty terms that have been discussed (and provided with the data) are: the C matrix= $C^{\nu\nu'}$, the beta matrix=β^p_k, the pixel weight=N_p, the JCJ gain term=J^ν, the P_gEP_g term=G^ν, the PUP temperature= uncertainty U, and the PTP absolute temperature uncertainty=T. Then the variance in the result A^a is a matrix V^{ab} and

$$V^{ab} = H^a_{p\nu} H^b_{p'\nu'} [C^{\nu\nu'}(\delta^{pp'}/N_p + \beta^p_k \beta_{p'k} + .04^2) + S^{p\nu} S^{p'\nu'}(J^\nu J^{\nu'} + G^\nu G^\nu \delta^{\nu\nu'}) + P^\nu P^{\nu'}(U^2 \delta^{pp'}/N_p + T^2)] \tag{35}$$

with some abuse of the summation convention. The $\delta^{pp'}$ is the familiar delta function which is one if the indecies match and zero otherwise and $P^\nu = \partial \mathrm{Planck}(2.728, \nu)/\partial T$.

After reading all of the error analysis and data devoted to error estimation one might come to the erroneous conclusion that the error estimation is difficult and complex. All of the processing has maintained the separability of the pixels and frequencies. This means that in operations involving fits (an average is a particular type of fit) in the pixel domain alone, only the $1/N_p$ and the β terms need to be treated. The $C^{\nu\nu'}$, G^ν, J^ν, and P^ν merely tag along and multiply the uncertainty of the result. Similarly when making a fit on the frequencies only (as in our line and temperature fits) only the $C^{\nu\nu'}$, G^ν, and J^ν need to be treated and the $1/N_p$ and the β terms are carried along and multiply the uncertainties of the results. If a fit involves both pixels and frequencies in a way that can be expressed as a product of operators then the pixels and frequencies can be treated separately.

Further simplification can be realized by noting that in any particular example only one or two terms may completely dominate the uncertainty. A few rules of thumb can be applied in determining the uncertainty terms that need to be used in analyzing the *FIRAS* data. Five examples of the use of the uncertainty terms in particular situations are provided. The rules of thumb are:

1. When only a few pixels are used (with less than 200 observations) and the signal is small, the detector noise alone is adequate.

2. When the signal levels are large, gain uncertainties (P_gEP_g, JCJ) dominate.

3. When data are averaged over large areas (more than ∼400 observations) the destriper uncertainties dominate.

4. When large scale structure in the spectra is analyzed JCJ is important.

5. When the absolute temperature is desired (in almost any data set) ONLY the PTP need be considered.

In any case when the final residual from any fit or model is made one should look over the list of FEF errors for one final sanity check.

These examples are given as questions to be asked of the data:

1. Question: What is the spectrum of the north Galactic pole?

   Dominant Uncertainty: C

   Explanation: In order to answer this question we select all of the pixels with Galactic latitudes greater than 87° and average them weighted by the number of pixels. We notice that the total number of observations which went into these spectra is 6. Since this number is less than 200, we expect that only the C uncertainty is important. The only part of the analysis where this is not true is for the absolute temperature of the CMBR.

2. Question: Is there H_2O absorption at the Galactic center?

   Dominant Uncertainty: $P_g E P_g$

   Explanation: Here we average together spectra within 5° of the center of the Galaxy. If the C uncertainty is the only uncertainty included in the analysis, it would appear that there is a very significant dip in the average spectra at the H_2O line. However, the dip is only a small fraction of the total power at the line. Because the signal level is high the gain fluctuations as a function of frequency (the $P_g E P_g$ gain) are important. Taking these uncertainties into account shows that the dip is a $\sim 3\ \sigma$ effect.

3. Question: What is the temperature of the Galactic dust?

   Dominant Uncertainty: JCJ

   Explanation: We make a weighted average of the Galaxy spectra. In this weighted average we note there are approximately 760 observations. This tells us that C is unlikely to be a good estimate of the uncertainty. The $P_g E P_g$ uncertainty is more important than the C, but the JCJ is still more important. Although the $P_g E P_g$ uncertainty is big at any given frequency, fitting the temperature over frequency reduces the effect of the $P_g E P_g$. Since the JCJ shows a systematic trend correlated from frequency to frequency, such an error reduction does not occur for the temperature fit. Consequently, the JCJ uncertainty dominates.

4. Question: What is the absolute temperature of the CMBR?

   Dominant Uncertainty: PTP

Explanation: The PTP is the largest error, but it is only important when estimating the absolute temperature. This uncertainty is derived from the uncertainty in the absolute temperature of the external calibration thermometers. Since the sky on average looks VERY much like the external calibrator, we only need to know what the external calibrator temperature is to get the CMBR temperature. But there is a 2 mK uncertainty in the calibrator temperature, hence there is a 2 mK error in the absolute temperature, even though there is enough signal-to-noise to see temperature variations of 10 μK.

5. Question: Can we see the temperature anisotropy that is seen by the *COBE DMR*?

   Dominant Uncertainties: PUP, D

   Explanation: The PUP uncertainty becomes important when looking for anisotropy in the CMBR. Since we are looking for variation in the absolute temperature across the sky, JCJ and PTP are unimportant. Fitting the CMBR temperature over frequency reduces the P_0EP_0 and C uncertainties. Thus the dominant uncertainty is C, but PUP is of the same order so it must be included too. This (and related questions) is the only type of modeling where the PUP uncertainty is important.

Note that in four of the five examples only one of the error terms is needed. And in Question 5 only two terms need to be considered. Further the full matrix was not needed as the error can be calculated after the final result was obtained.