The Astrophysical Journal, 491:409–413, 1997 December 10

SEARCH FOR PERIODICITIES IN THE HOMESTAKE SOLAR NEUTRINO DATA

P. A. Sturrock,[1] G. Walther,[2] and M. S. Wheatland[1]
Received 1997 March 18; accepted 1997 July 16

ABSTRACT

We evaluate a χ^2 statistic to test against the Homestake data the hypothesis that the neutrino flux from the Sun is constant. We use estimates of standard deviations derived 1000 simulations of the sequence of 108 runs, and we also use two procedures for deriving proxies for the standard deviation from the experimental data. All tests indicate that the hypothesis should be rejected; the significance level ranges from 5.8% to 0.1%.

We also search for evidence of periodicities in the neutrino flux by evaluating the log likelihood of finding the actual count rates in a model in which the neutrino flux is modulated with a sinusoidal term. We consider a range of values of the frequency (0–20 cycles yr^{-1}) and, for each frequency, adjust the modulation parameters to maximize the likelihood. We find no evidence of modulation at the frequency of the solar cycle. A 1000 shuffle test and 1000 simulations using error estimates taken from the simulations yield no evidence for either the quasi-biennial (2.2 yr) periodicity or the Rieger (157 day) periodicity. However, simulations based on the experimental error estimates yield significance levels of 1% and 2.7% for the quasi-biennial periodicity, and 2% and 0.2% for the Rieger periodicity.

We have also looked for evidence of modulation at a frequency that might be related to the solar rotation frequency. We have adopted a search band of 12.4–13.1 cycles yr^{-1}, corresponding to the 1 year lower sideband (synodic frequency) of the rotation frequency of the Sun's radiative zone, as estimated from helioseismology. There is indeed a peak in that band, at 12.88 cycles yr^{-1}, that according to the simulation test is significant at the 3% level. However, we also find evidence of four sidebands near 10.88, 11.88, 13.88, and 14.88 cycles yr^{-1} that may be due to the departure of the rotational axis from the normal to the ecliptic. We introduce a correlation measure formed from the powers at a "fundamental" and at four sidebands. None of 1000 shuffle tests, and only one of 1000 simulations, yield values of the correlation measure as large as that formed from the experimental data. These tests offer support, at the 0.1% and 0.2% significance level, respectively, for the proposition that the neutrino flux is modulated at a frequency that could be the synodic frequency corresponding to a sidereal rotational frequency of 13.88 cycles yr^{-1} (440 nHz) of the Sun's radiative zone.

Subject headings: elementary particles — methods: statistical — Sun: interior — Sun: particle emission

1. INTRODUCTION

The Homestake neutrino measurements, comprising 108 runs extending from 1970.281 to 1994.388, provide the longest sequence of data giving information about the deep solar interior (see, e.g., Davis & Cox 1991). From the very beginning, there has been keen interest in the possibility that these data yield evidence for real variations in the neutrino flux. This question has been addressed either by searching for correlation between the Homestake measurements and an indicator of solar activity such as the sunspot number (see, e.g., Bahcall, Field, & Press 1987; Bahcall & Press 1991; Bieber et al. 1990), or by searching for periodicity in the data (see, for instance, Haubolt & Gerth 1990). These important issues concerning the solar neutrino flux have recently been authoritatively reviewed by Davis (1995).

Variation in the neutrino flux can in principle result from either or both of two causes: (1) "fluctuation" (the nuclear burning may be time variable, as recently suggested by Grandpierre 1996), or (2) "modulation" (the solar neutrino flux may be modulated either by the solar magnetic field, as suggested by Voloshin, Vysotskii, & Okun 1986a, 1986b, or by some other process).

In view of the fact that the Homestake neutrino data have recently been completely reanalyzed by the Homestake team, resulting in changes in the uncertainties of the measurements, it seems timely to examine these questions once more. In § 2 we present a statistical test for the constancy of the solar neutrino flux, as measured by the Homestake experiment, and find evidence that the flux is not constant. In § 3 we report spectrum analysis of the Homestake data and look for evidence of three well-known "low-frequency" solar periodicities: the 11 year solar cycle, the 780 day quasi-biennial periodicity, and the Rieger 157 day periodicity. We find no evidence of the solar cycle, but some tests show some evidence for the quasi-biennial periodicity and for the Rieger periodicity.

If the solar neutrino flux is modulated by the solar magnetic field, the form of the resulting time variation will depend upon the location of that field. Any magnetic field in the convection zone is unlikely to have a long-lasting longitudinal structure, due to the strong differential rotation, so it is unlikely to produce any "high-Q" rotational modulation, although it could lead to modulation with the period of the sunspot cycle. On the other hand, a "fossil field" in the radiative zone (that may offer an alternative explanation of the solar cycle [see, e.g., Sturrock 1997]) is necessarily

[1] Center for Space Science and Astrophysics, Stanford University, Stanford, CA 95305-4060.
[2] Statistics Department, Stanford University, Stanford, CA 95305-4065.

long-lived and may therefore produce a high-Q rotational modulation of the neutrino flux. For this reason, it seems important to search for evidence of a high-Q and "high-frequency" periodicity that may be related to the internal solar rotation of the Sun. Such a search is carried out in § 4.

The results of our analysis are discussed briefly in § 5.

2. TEST FOR TIME VARIATION

We can in principle determine whether the neutrino flux is time-varying by determining whether the scatter in measurements is significantly larger than we would expect if the flux were constant. B. Cleveland of the Homestake team has developed a code that can be used both to simulate the Homestake experiment and to analyze either data acquired by the actual experiment or data generated by the simulation (B. Cleveland 1996, private communication). Cleveland has generously provided us with a copy of this code, and Kenneth Lande has generously provided us with the results of their complete reanalysis of the Homestake data (K. Lande 1996, private communication).

We have used this code to generate 1000 simulations of the actual sequence of 108 runs, based on an assumed constant Ar production rate of 0.475 atoms day^{-1}, the value obtained by the Homestake team by their maximum likelihood analysis of the actual sequence of 108 runs. This code begins by simulating (for a given flux of neutrinos) the Poisson process that governs the creation of radioactive Ar atoms due to the conversion of Cl atoms. The code then simulates the known background radiation and the Poisson process of the decay of the Ar atoms, so producing a series of times at which detection events would be registered in the counter following the Ar extraction operation. The simulations generated by this code mimic the real experiment as accurately as possible. The exposure time, the experimental efficiencies of extracting and counting, the length of counting, the counter resolution, and the background radiation in the counter have all been chosen to be identical to those of each real run.

For each run of each simulation, the series of detection times was then analyzed with exactly the same maximum likelihood program that had been used to determine the production rate and the confidence limits for each run in the Homestake experiment. Each simulation ($\alpha = 1, \ldots 1000$) therefore yields 108 estimates $g_{i,s\alpha}$ ($i = 1, \ldots 108$) of the production rate g_i, and similarly for f_i and h_i, the lower and upper 68% confidence limits, all measured in ^{37}Ar atoms day^{-1}. We found that, for each run, the 1000 estimates of the production rate may be fitted approximately (but only approximately) to a Gaussian distribution. Hence, for each run, we could determine from the simulations a standard deviation $\sigma_{i,s}$. The departure from a Gaussian form is not crucial for the analysis in this section.

For simulation α, we may now form the χ^2 statistic.

$$\Gamma_{s\alpha} = \sum_i \left(\frac{g_{i,s\alpha} - \bar{g}_{s\alpha}}{\sigma_{i,s}}\right)^2 , \tag{1}$$

where $\bar{g}_{s\alpha}$ is chosen to minimize the statistic. We also compute the same statistic from the experimental data,

$$\Gamma_e = \sum_i \left(\frac{g_{i,e} - \bar{g}_e}{\sigma_{i,s}}\right)^2 , \tag{2}$$

where $\bar{g}_e$ is chosen to minimize this statistic. We find that $\Gamma_e = 133.58$. We find that 57 of the 1000 simulations give values of $\Gamma_{s\alpha}$ larger than Γ_e, contradicting the "null hypothesis" that the neutrino flux is constant at a significance level of 5.8%. For comparison, it is interesting to note that, for 107 degrees of freedom (dof), this value of the χ^2 statistic has a significance of 4.2%. The two estimates are quite close, even though the distribution of estimated values is not exactly Gaussian.

Bahcall & Press (1991) attached less significance to runs earlier than run 49, for reasons set out in their article. If we consider a reduced data set, namely, the 81 runs beginning at run 49, we find that 35 of 1000 simulations result in a value of Γ_α larger than Γ_e, giving a significance level of 3.6%. When compared to a χ^2 distribution with 80 dof, we obtain a significance level of about 1.8%, so that the two estimates are once again similar.

We have also considered two "empirical" error estimates, which may be regarded as proxies for the standard deviation. We denote these by $\sigma_{i,ea}$ and $\sigma_{i,em}$, defined as follows,

$$\sigma_{i,ea} = \tfrac{1}{2}(h_{i,e} - f_{i,e}) \tag{3}$$

and

$$\sigma_{i,m} = \max\,[(g_{i,e} - f_{i,e}), (h_{i,e} - g_{i,e})] , \tag{4}$$

where the subscripts ea denote "empirical, average," and the subscripts em denote "empirical, maximum." If we use the proxy ea (which was used by Bahcall et al. 1987), the experimental data yield the value $\Gamma_{ea} = 219.01$, and we find that none of the 1000 simulations (with the $\sigma_{i,ea}$ given by the simulation) gives a larger value of the statistic, implying that the assumption of constant generation rate may be rejected at a significance level of 0.1%. In considering the "reduced" data set, we find the same result. If we use the proxy em (which was used by Bieber et al. 1990), the experimental data yield the value $\Gamma_{em} = 124.65$, and we find that none of the 1000 simulations gives a larger value of the statistic, implying that the assumption of constant generation rate may be rejected at a significance level of 0.1%. The same result is obtained also for the reduced data set.

It is appropriate to inquire into the reason that two different uncertainty estimates lead to strikingly different significance estimates. We note (as has been pointed out by Bieber et al. 1990) that the empirical error estimates produced by the maximum likelihood procedure tend to be correlated with the estimated production rate, in that smaller error estimates are associated with smaller estimates of the production rate. The reason for this appears to be that the error estimates are obtained by integrating the likelihood function using a prior distribution that is restricted to the positive half-line; a small estimate of the production rate normally arises from a likelihood function that is concentrated near zero, but this also leads to a smaller error estimate. On the other hand, the frequentist error estimates obtained from the distribution of estimates generated by the simulations do not lead to such a correlation between production-rate estimates and error estimates; the error estimates derived from the simulations are determined by the various parameters of the experiment (such as exposure time).

In this context it is also interesting to note that the actual production rate estimates made by the Homestake team exhibit greater variability (i.e., have more large and more small values) than do the simulated data. The χ^2 statistic (eq. [2]) is strongly affected by this difference in variability, since a small error estimate leads to a small term in the denominator of the corresponding term of the χ^2 formula. It appears, therefore, that the difference in estimated significance levels may be due in part to time variation of the Homestake data, as the test that uses the error estimates produced by the maximum likelihood procedure is more sensitive to time variation.

Although different procedures yield different error estimates, we may use any such set of error estimates, in conjunction with a Monte Carlo procedure, to arrive at a significance estimate. If two different estimates of the significance level of a test are both valid, the less stringent result is clearly compatible with the more stringent result, but the converse is not true. We therefore conclude that the null hypothesis that the neutrino flux is constant may be rejected at the 0.1% confidence level.

We have also repeated the analysis of this section assuming production rates of 0.45 atoms day^{-1} and 0.50 atoms day^{-1} for the simulations. The results are substantially the same as those we obtained with 0.475 atoms day^{-1}, showing that the results are not sensitive to the assumed production rate.

3. LOW-FREQUENCY SPECTRUM ANALYSIS

One may use spectrum analysis as an attempt to characterize a time-varying series. If any periodicity is discovered, one may then seek to relate that periodicity to a known *solar* periodic process. The usual procedure for assessing the strength of the case that an apparent periodicity is real is to determine the probability that it could appear by chance.

We have carried out a spectrum analysis of the Homestake neutrino data over a wide range of frequencies, which we measure in cycles per year unless otherwise specified. However, since we are searching for evidence of specific frequencies, it is convenient to discuss our results in two sections dealing (in the next section) with high-frequency periodicities (comparable to the Sun's rotation frequency) and (in this section) with low-frequency periodicities. In each section, we first look for prima facie evidence that a suspected periodicity is present, and if it appears to be present, we then proceed to assess the significance of that periodicity.

In this section we search for three known low-frequency solar periodicities. One of these is the familiar solar cycle with a period of about 11 years, or, equivalently, a frequency close to 0.09 cycles yr^{-1}. The second is the quasi-biennial periodicity (see, e.g., Nesme-Ribes et al. 1993), which has a period in the range 730–810 days (frequency = 0.45–0.50 cycles yr^{-1}). (Sakurai 1979, 1981 has previously presented evidence that this periodicity is present in the Homestake data.) The third is the "Rieger" periodicity (Rieger et al. 1984; Bai & Sturrock 1987), which has a period in the range 152–159 days (frequency = 2.30–2.40 cycles yr^{-1}).

The Homestake data comprise (for each run) the "begin" date $t_{b,i}$, the "finish" date $t_{f,i}$, and the estimates $f_{i,e}$, $g_{i,e}$, and $h_{i,e}$ used in § 2. These data are so sparse that it is not possible to generate a meaningful estimate of the spectrum of the neutrino generation rate by a simple Fourier transform procedure. We therefore use the maximum likelihood estimation procedure (see, for instance, Bevington & Robinson 1992).

The log likelihood Λ of a series of events is defined as

$$\Lambda = \ln L = \sum_i \ln p_i \,, \tag{5}$$

where p_i represent the probabilities of the events under consideration. We estimate the likelihood of the actual events occurring in a model in which the ^{37}Ar generation rate, $\gamma(t)$, has the form

$$\gamma_m(t) = C + A \cos(2\pi\nu t) + B \sin(2\pi\nu t) \,. \tag{6}$$

The expected value of the count rate for each run may be found from

$$g_{i,m} = \frac{\int_{t_{b,i}}^{t_{f,i}} \gamma_m(t) e^{-\lambda(t_{f,i}-t)} dt}{\lambda^{-1}(1 - e^{-\lambda(t_{f,i}-t_{b,i})})} \,, \tag{7}$$

since ^{37}Ar atoms decay (with a time constant λ of 0.0198 day^{-1}, corresponding to a half-life of 35 days), and Homestake estimates of the count rate are based on the assumption that the neutrino flux is constant during the run. The probabilities p_i are therefore expressible as

$$p_i = P_i(g_{i,e} | A, B, C, \nu) \,. \tag{8}$$

For given experimental data and for each value of the frequency ν, we choose A, B, and C to maximize the likelihood, subject, however, to the restriction that $A^2 + B^2 \leq C^2$ to ensure that the generation rate is nonnegative. We then determine the likelihood Λ that the data would be produced by those values of A, B, and C (more precisely, we determined the increase in log likelihood from the best fit with $A = B = 0$).

As in § 2, we represent the probability distribution in equation (8) as a Gaussian distribution with standard deviation $\sigma_{i,s}$, so that equation (5) may be expressed as

$$\Lambda = -\frac{1}{2} \sum_i \left(\frac{g_{i,e} - g_{i,m}}{\sigma_{i,s}} \right)^2 \,. \tag{9}$$

To facilitate the comparison of different log likelihood estimates, it has been convenient to normalize them so that each has a mean value of unity. For convenience, we refer to the normalized log likelihood as the "power." The normalization was carried out for the 2000 log likelihood estimates computed over the frequency range 0–20 cycles yr^{-1} in steps of 0.01 cycles yr^{-1}. The precise definition of the likelihood used in these calculations is not crucial, since we do not propose to read off significance estimates directly from the heights of the peaks in the spectrum.

In examining the spectrum in the neighborhood of the frequency of the solar cycle (0.09 cycles yr^{-1}), we find no evidence of a peak. This is compatible with the analysis of Haubolt & Gerth (1990) and with the recent result that there is no evidence of the solar cycle in the Kamiokande data (Fukuda et al. 1996), but appears to be incompatible with the analysis of Bahcall & Press (1991), who found that the neutrino flux is anticorrelated with solar activity.

In examining the spectrum in the search band of the quasi-biennial periodicity, we find that there is indeed a peak with power 2.58 at frequency 0.49, corresponding to a period of about 745 days.

We have adopted two procedures for obtaining distribution-free significance estimates. One is the "shuffle test" that was used by Bahcall & Press (1991) in their analysis of the Homestake data. We form many "pseudosequences" by randomly changing the order of the runs (retaining for each run the duration, the dead time if any up to the beginning of the next run, the estimated count rate, and the error estimates), and then compute the spectrum for each pseudosequence formed in this way. (Note that we did not ensure nonnegativity in the shuffle test, as that would have been prohibitively demanding on our computer resources. However, we did carry out a short run of the shuffle test in which nonnegativity was ensured; the results differed little from a similar run that did not ensure nonnegativity.) We then determine the fraction of the total number of this hypothetical set of pseudosequences in which the power (unnormalized) exceeds that of the actual sequence in the search band 0.45–0.50 cycles yr^{-1}. This fraction is found to be 13.1%, so that, according to this analysis, the apparent peak at the quasi-biennial frequency is not statistically significant.

The other procedure is to determine the fraction of the 1000 simulations for which the power in the search band exceeds the power of the peak in the actual spectrum. We find that, for the quasi-biennial periodicity, 133 of the 1000 simulations yielded a power larger than that of the actual spectrum, yielding a significance level of 13.4%.

We have applied the same procedure to the Rieger periodicity. The maximum power is found to be 3.11 at $\nu = 2.32$, corresponding to a period of about 157 days. The shuffle test, with a search band of 2.30–2.40 cycles yr^{-1}, indicates that this peak is not significant. From the simulations, it appears to be significant at the 10% level.

Since the distribution of count-rate estimates from the simulations is not exactly Gaussian in form, we have also repeated these calculations with the assumption that the estimates are better represented by a Cauchy distribution. This is known to be in general a more robust procedure. When this procedure is applied to the quasi-biennial periodicity, we do not find the peak to be significant. When it is applied to the Rieger periodicity, the significance level improves to 3.5%.

4. HIGH-FREQUENCY SPECTRUM ANALYSIS

We now examine the possibility that the solar neutrino flux may be modulated by inhomogeneities within the Sun, as discussed in § 1. If the axis of rotation of the inhomogeneities were normal to the ecliptic, and if the fundamental (sidereal) rotational frequency were ν_R, we could expect to detect the lower sideband or "synodic" frequency $\nu_R - 1$ and its harmonics. We find from simulations that frequencies of order 13 cycles yr^{-1} are difficult but not impossible to extract from the Homestake data, but any harmonics would be virtually impossible to extract. If the axis of rotation were to differ significantly from the normal to the ecliptic, we might also detect the fundamental frequency and possibly also the upper sideband $\nu_R + 1$. If the inhomogeneity were to be localized to a small latitude band, there would be a seasonal variation in the modulation: this would lead to the appearance of the synodic frequency $\nu_R - 1$ and to sidebands of that frequency, such as $\nu_R - 1 \pm 1$ and possibly $\nu_R - 1 \pm 2$, etc. (See, e.g., Sturrock & Bai 1992.)

We argued in § 1 that a long-lived magnetic structure, which could lead to a high-Q modulation of the neutrino flux, is more likely to be found in the radiative zone than in the convection zone, since the latter exhibits strong differential rotation whereas the former is more nearly in rigid rotation. The *SOHO* MDI team (Kosovichev et al. 1997; see also Lang 1997) has recently derived an estimate of the internal rotation-rate profile of the Sun from the first few months of data. In the latitude band -30 to $+30$ and for normalized radius 0.4–0.7, the estimated rotation rate is in the band 425–435 nHz. A recent analysis of Global Oscillation Network Group (GONG) data (Thompson et al. 1996) yields estimates in the band 430–440 nHz. Based on these estimates, we search for evidence of rotational modulation corresponding to a fundamental frequency ν_R in the range 13.4–14.1 cycles yr^{-1}, which corresponds approximately to 425–445 nHz. We would expect the dominant component to have a frequency in the range 12.4–13.1 cycles yr^{-1}, and we therefore adopt this as our search band.

Figure 1 shows a plot of the power in the range 10–15. We see that the biggest peak in this range, namely, a peak with $S = 4.21$ at $\nu = 12.88$, falls within the search band. According to the shuffle test, there is a probability of 11.3% of finding such a peak in the search band by chance, and analysis of the simulations yields a corresponding probability of 3%. In view of the latter estimate, the peak seems interesting.

From inspection of Figure 1, we find evidence not only for a peak near 12.9 but also for peaks near 10.9, 11.9, 13.9, and 14.9. Specifically, we find the following peaks: $S = 2.58$ at $\nu = 10.83$, $S = 3.43$ at $\nu = 11.85$, $S = 3.79$ at $\nu = 13.85$, and $S = 2.13$ at $\nu = 14.88$. As we noted above, such sidebands can occur if the rotation axis is not normal to the ecliptic. We have therefore sought to assess the probability that such a group of five peaks, with spacing approximately $\Delta\nu = 1$, might occur by chance. This calls for something like a correlation analysis.

We have formed a "correlation index" $\Gamma_5(\nu)$ by multiplying the powers at five frequencies with spacing $\Delta\nu = 1$:

$$\Gamma_5(\nu) = S(\nu - 2)S(\nu - 1)S(\nu)S(\nu + 1)S(\nu + 2) . \quad (10)$$

The result is shown in Figure 2. We see that there are three notable peaks of 68.20 at 11.85, 96.90 at 12.85, and 95.35 at 13.85. In 1000 shuffles, we find in the prescribed search band no peaks larger than the peak at 12.85. This indicates that the null hypothesis, that there is no periodicity with related

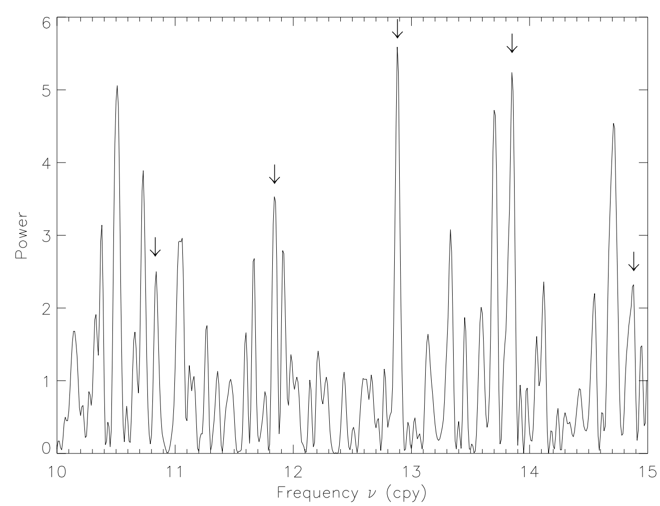

FIG. 1.—Spectrum of the Homestake neutrino data over the band $\nu = 10$–15, obtained by maximum likelihood analysis based on a Gaussian error distribution, the standard deviations being taken from the results of simulations. Arrows indicate the peaks at 10.83, 11.85, 12.88, 13.85, and 14.88.

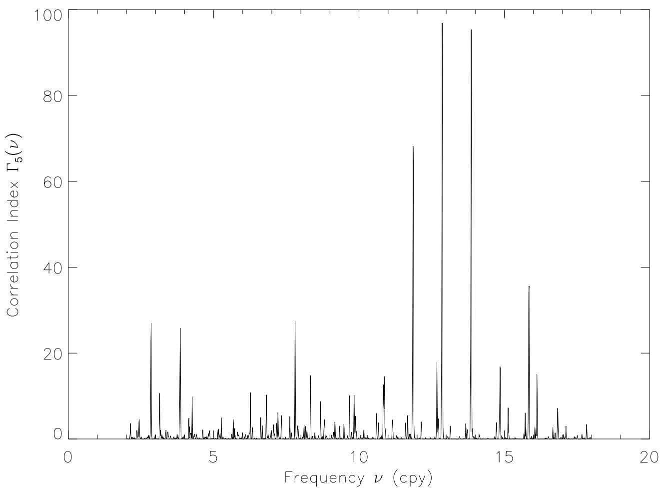

FIG. 2.—Correlation index, defined by eq. (8), over the band $\nu = 2$–18

sidebands in this search band, may be rejected at the 0.1% significance level. We have also examined 1000 simulations and found one peak of the correlation index within the search band larger than that found in the data. This test indicates that the null hypothesis, that there is no periodicity with related sidebands in this search band, may be rejected at the 0.2% significance level.

We have considered several variants of this procedure. We have examined correlation indexes formed from two, three, and seven terms, as well as from five. The apparent significance of the peak at 12.88 increases progressively with the order of the correlation index. We have also replaced the power by related variables such as $(1 + S)/2$ or a non-parametric proxy for the power derived from the rank orders of the peaks in the spectrum. Examination of spectra formed in this way shows that the quintuplet of peaks shows up in each case.

5. DISCUSSION

The analysis of § 2 confirms the analysis of other investigators, which is that the data yielded by the Homestake neutrino experiment are not consistent with a steady neutrino flux. The analysis of § 3 indicates that the neutrino flux does not vary with the solar cycle, agreeing with some previous claims and disagreeing with others.

The analysis of § 3, based on error estimates taken from the simulations, offers only a suggestion that the flux may vary with the quasi-biennial or Rieger periodicities. However, we have also analyzed the data on the basis of the error estimates *ea* and *em* introduced in § 2 (see eq. [3] and [4]). These give very different results. In examining the quasi-biennial periodicity on the basis of *ea*, the shuffle test yields a significance estimate of 8.1% but simulations yield an estimate of 1.0%. On using *em*, we obtain an estimate of 11.4% from the shuffle test and 2.7% from the simulations. In examining the Rieger periodicity on the basis of *ea*, the shuffle test indicates that the peak is not significant (38.8%), but simulations yield an estimate of 2.0%. On using *em*, we obtain an estimate of 8.6% from the shuffle test and 0.2% from the simulations. Unless some of these tests are invalid, it seems likely that the neutrino flux varies with the periods of the quasi-biennial periodicity and the Rieger periodicity.

Our search for periodicities in the Homestake data yields strongest evidence for a periodicity at about $\nu = 12.88$. If real, this could be the synodic value of the rotation rate of the Sun's radiative zone. The case for this periodicity alone is modest, but the case for this periodicity plus its four nearest sidebands seems quite strong.

If the solar neutrino flux really is modulated at a frequency corresponding to the rotation rate of the radiative zone, it must be that some inhomogeneity in the solar interior is influencing the electron neutrino flux. This modulation could in principle be a result of either density variations (Mikheyev & Smirnov 1985; Wolfenstein 1978) or magnetic field variations (Voloshin et al. 1986a, 1986b), but the latter would seem to be the more promising explanation, since the Sun can accommodate much greater inhomogeneities in its magnetic field than in its density. As is well known (see, e.g., Bahcall et al. 1996), flavor oscillation may also explain the discrepancy between the expected neutrino flux and the measured flux.

For further information about the statistical analysis reported in this article, the interested reader is invited to contact G. W. (walther@playfair.stanford.edu).

We are grateful to many colleagues for helpful discussion of this work. In particular, we wish to thank Bruce Cleveland, Ray Davis, and Kenneth Lande for providing abundant information concerning the Homestake experiment, and William Press, who served as referee of the first version of this article and made many comments and suggestions that led us to completely rework our analysis. This work was supported by NASA grants NAS 8-37334 and NAGW-2265 and by Air Force grant F49620-95-1-0008.

REFERENCES

Bahcall, J. N., Calaprice, F., McDonald, A. B., & Totsuka, Y. 1996, Phys. Today, July issue, 30
Bahcall, J. N., Field, G. B., & Press, W. H. 1987, ApJ, 320, L69
Bahcall, J. N., & Press, W. H. 1991, ApJ, 370, 730
Bai, T., & Sturrock, P. A. 1987, Nature, 327, 601
Bevington, P. R., & Robinson, D. K. 1992, Data Reduction and Error Analysis for the Physical Sciences (2d ed.; New York: McGraw-Hill)
Bieber, J. W., Seckel, D., Stanev, T., & Steigman, G. 1990, Nature, 348, 407
Davis, R., Jr. 1995, TAUP '95 Conference Proceedings (in press)
Davis, R., Jr., & Cox, A. N. 1991, Solar Interior and Atmosphere, ed. A. N. Cox, W. C. Livingston, & M. S. Matthews (Tucson: Univ. Arizona Press), 51
Fukuda, Y., et al. 1996, Phys. Rev. Lett., 77, 1683
Grandpierre, A. 1996, A&A, 308, 199
Haubolt, H. J., & Gerth, E. 1990, Sol. Phys., 127, 347
Kosovichev, A. G., et al. 1997, Sol. Phys., 170, 43
Lang, K. R. 1997, Sci. Am., March issue, 40
Mikheyev, S., & Smirnov, A. 1985, Soviet J. Nucl. Phys., 42, 913
Nesme-Ribes, E., Sokoloff, D., Ribes, J. C., & Kremiovsky, M. 1993, The Solar Engine and Its Influence on Terrestrial Atmosphere and Climate, ed. E. Nesme-Ribes (Berlin: Springer), 71
Rieger, E., et al. 1984, Nature, 312, 623
Sakurai, K. 1979, Nature, 278, 146
———. 1981, Sol. Phys., 74, 35
Sturrock, P. A. 1997, Sol. Phys., in press
Sturrock, P. A., & Bai, T. 1992, ApJ, 397, 337
Thompson, M. J., et al. 1996, preprint
Voloshin, M. B., Vysotskii, M. I., & Okun, L. B. 1986a, Soviet J. Nucl. Phys., 44, 440
———. 1986b, Soviet Phys.—JETP, 64, 446
Wolfenstein, L. 1978, Phys. Rev. D, 17, 2369