JOURNAL OF MAGNETIC RESONANCE 98, 483–500 (1992)

Comparison of Fourier and Bayesian Analysis of NMR Signals. I. Well-Separated Resonances (The Single-Frequency Case)

JOHN J. KOTYK,* NORMAN G. HOFFMAN, AND WILLIAM C. HUTTON

Monsanto Company, 700 Chesterfield Parkway North, St. Louis, Missouri 63198

AND

G. LARRY BRETTHORST AND JOSEPH J. H. ACKERMAN

Department of Chemistry, Washington University, 1 Brookings Drive, St. Louis, Missouri 63130

Received June 17, 1991; revised November 13, 1991

Analysis of nuclear magnetic resonance data requires an estimation of the parameters (e.g., frequencies and amplitudes) that characterize the NMR free-induction-decay data. Bayesian probability theory provides a rigorous formalism for optimal parameter estimation and use of prior information. Although specific algorithms for time-efficient implementation of Bayesian methods have been presented (G. L. Bretthorst, *J. Magn. Reson.* **88**, 533, 552, 571, 1990; G. L. Bretthorst, *J. Magn. Reson.*, in press), acceptance of such methods requires demonstration of substantial improvements in the accuracy of parameter estimates. Toward this end, we report a comparison between the discrete Fourier transform (DFT) method and Bayesian analysis for estimating the signal *frequency* and *amplitude* of a single-frequency NMR resonance as a function of the signal-to-noise (S/N) ratio of the FID data. The results and methods are also applicable to data composed of multiple, well-separated frequency components. Parameter estimates are made both with and without prior knowledge of the decay rate and/or phase of the NMR signal. Under conditions where prior information about the signal phase and decay rate constant is not available, Bayesian analysis provides more precise estimates of the signal frequency and continues to do so considerably after the DFT method fails due to poor S/N levels. In accordance with theory, the Bayesian and DFT methods yield identical *frequency* estimates when the DFT estimates are obtained from a zero-padded absorption spectrum when prior information about both the decay rate constant (i.e., matched exponential filter) and the signal phase is available. In all cases, Bayesian analysis is substantially more precise than the DFT method for estimating signal *amplitudes*. Reasons for the differences observed between the two analysis techniques are discussed in detail. At this level of validation, Bayesian analysis offers distinct advantages over DFT procedures for NMR parameter estimation.

Two major nuclear magnetic resonance parameters that a spectroscopist often wishes to determine from time-domain NMR data are the frequencies and the amplitudes of the observed signals. Accurate estimation of these quantities is key to elucidating

* To whom correspondence should be addressed.

0022-2364/92 $5.00

structures and identifying dynamic processes on a molecular level. Most NMR laboratories currently estimate these parameters by examining a frequency-domain NMR absorption spectrum that results from a discrete Fourier transform (DFT) (*1–4*) of the complex time-domain free-induction-decay NMR data. The DFT is usually performed in conjunction with zero-filling and appropriate apodization (*5, 6*) of the time-domain data to help improve the frequency-domain digital resolution and signal-to-noise (S/N) ratio, respectively. Frequency estimates for NMR signals are then obtained using secondary procedures, ranging from simple peak-picking routines to nonlinear least-squares fitting procedures, applied to the real part of the DFT (i.e., the absorption spectrum). Similarly, integrated intensities for a defined region of the spectrum spanning the peak of interest provide quantitative estimates of the concentration of spins giving rise to the NMR signal.

Although NMR data analysis using the DFT is clearly useful, a number of well-documented deficiencies of the method are known to exist (*7, 8*). For example, large errors in the NMR parameter estimates can result for FID data containing multiple decaying sinusoidal frequencies which overlap in the corresponding frequency-domain absorption spectrum. In such cases, peak maxima in the frequency spectrum do not necessarily correspond with the actual frequencies, and separate signal amplitudes cannot be determined via integrated areas. In addition, no means exist to determine the quality or precision of a single independent measurement for parameter estimates obtained from the DFT absorption spectrum or from nonlinear least-squares fitting procedures. Baseline distortions in the absorption spectrum also limit DFT performance, especially when rapidly decaying signals are present in the FID data. Such distortions result from finite pulse widths and delays prior to data sampling and are one of the major sources of baseplane distortions, including t_1 and t_2 ridge artifacts (*7*), that hinder the analysis of two-dimensional and other multidimensional NMR experiments. NMR FID data analysis using the DFT also suffers at low frequency-domain S/N levels, where manual and automatic data phasing become difficult, making it impossible to differentiate between large random noise peaks and incorrectly phased low-intensity resonances.

Recently, non-Fourier-based alternatives have been used to examine time-domain NMR FID data. For such methods to be accepted, their ability to obtain precise and accurate estimates of NMR parameters must be demonstrated. Two such methods, maximum entropy (MEM) and linear prediction (LP), have shown considerable promise in providing reliable NMR parameters estimates (*8*). Reports that these methods offer simultaneous improvements in S/N and resolution (*9–11*) have been made, but these claims are controversial (*12, 13*). MEM and LP have also been evaluated for use with *in vivo* NMR spectroscopy (*14*). As potential solutions to some of the problems associated with the DFT, new analysis methods must be directly compared to the DFT procedures. Such a comparison has been reported for MEM (*15, 16*).

Another non-Fourier analysis technique, Bayesian probability theory, has also been applied to the problem of NMR parameter estimation (*17–22*). Bayesian analysis permits the introduction of "prior information" about the NMR signal and, by rigorous use of probability theory (*23*) and Bayes' theorem (*24*), selects the model that is most consistent with this information and the time-domain NMR data. As a mature, well-understood discipline, NMR spectroscopy offers an ideal area for application of Bayes-

ian analysis. Good mathematical models for NMR data are available and can be used to improve the estimates of the frequencies and amplitudes of sinusoidal NMR signals. In addition, uninteresting or nuisance parameters are removed through the well-established process of marginalization. For example, in estimating the frequency, the signal amplitude and the decay rate constant can be treated as a nuisance parameter, while frequency and decay rate constants are nuisance parameters when estimating the signal amplitude. Like other analytical procedures, the ability of Bayesian methods to estimate NMR parameters from time-domain data is limited by the S/N ratio of the data. Bayesian methods do not improve the inherent S/N level of the experiment. Optimal parameter estimates can be extracted from the time-domain data by asking very specific questions about the data and then calculating a numerical probability for the answers to these questions. In practice, an estimate is chosen as the maximum in the posterior probability distribution that is calculated from a mathematical formulation of the question. Improved precision and accuracy for the parameter estimates are obtained by more efficient use of the prior information about the exact (or theoretical) nature of the data and the manner in which it has been acquired.

The present work reports an empirical comparison between the abilities of the DFT and Bayesian probability procedures to analyze well-separated resonance frequencies in NMR FID data. Our emphasis is on evaluating the precision that the two methods provide in estimating the frequency and signal amplitude of the NMR resonance under conditions of decreasing S/N levels. For brevity, we have excluded estimating the decay rate constants. Probability theory indicates that for n well-separated frequencies, optimal frequency estimates are obtained by selecting the n most probable single-frequency estimates. Therefore, to avoid redundancy, we have used a single-frequency model system. It can be formally shown (*21*) that in the absence of prior signal phase information, frequency estimates measured from a zero-padded, exponentially apodized (using a matched-weighting filter) DFT *power spectrum* are identical to those obtained using Bayesian analysis, if the NMR resonances are well separated. That is, the *power spectrum* is the optimal statistic for obtaining frequency estimates in the absence of prior phase information. If prior phase information is known, the *absorption spectrum* provides the optimal statistic for frequency estimation under the same conditions. Hence, we have made comparisons between the two analytical procedures in the presence and in the absence of prior knowledge of the signal phase and decay rate constant. In addition, estimates of the signal amplitude are reported by Bayesian procedures as the peak of the amplitude probability distribution and by DFT procedures as the integrated area of the NMR resonance of the absorption spectrum. Under ideal conditions, these two quantities reflect equivalent quantitative measures. Thus, we refer to both as the NMR *signal amplitude.*

EXPERIMENTAL METHODS

All ^1H NMR data were acquired on a Varian Unity 500 MHz NMR spectrometer equipped with a Varian $^1H/^{19}F$ NMR probe. A ^1H FID was collected for each of four D_2O samples containing differing amounts of $GdCl_3$ to shorten the decay rate of the observed residual HOD resonance. Linewidths for these samples were determined from the absorption spectrum to be 8, 32, 240, and 832 Hz following a DFT. All data were collected on nonspinning samples under conditions that avoid radiation damping

using identical acquisition parameters (bandwidth, 20 kHz; acquisition time, 1.237 s $\geq 10T_2^*$; double-precision complex data points, 24,736; recycle delay, $>20T_1$; temperature, 30°C). For data acquisitions using multiples of four phase-cycled transients, a time-domain RMS S/N ratio of at least 4000 was achieved for each sample FID. In all cases, the signal was placed −1000 Hz off resonance and had decayed to well below the noise level by 16K data points into the data set. NMR signal amplitudes were arbitrarily scaled to a value of 100 for convenience of display and discussion.

Each of the four FIDs was examined using the DFT and Bayesian analysis methods described below at noise levels characterized by a known standard deviation (σ). At each integer value of σ ranging from 1 to 50, 50 different individual Gaussian white noise sets were generated (*25*). This series of 10,000 (i.e., 4 FIDs; 2500 noise samples/FID) unique *combined data sets* (i.e., FID + Gaussian noise) was then analyzed to obtain frequency and signal amplitude estimates. Standard deviations for the NMR parameter estimates were computed at each σ from the 50 independent parameter estimates made for each of the combined data sets. Before deciding to use computer-generated noise, we verified that noise collected on the spectrometer gave statistically identical results when the audiofilter bandwidth was set to twice the spectral window. For the series of combined data sets, a σ value of 1 corresponds to a *time-domain* peak to RMS S/N of 100 and a σ value of 100 corresponds to a *time-domain* S/N of 1. On the basis of the results of the initial survey, some data series were extended to higher σ values for presentation purposes, as indicated in the text.

The Fourier number for each analysis procedure was adjusted to adequately define the peak shape and to help minimize computing time. Fourier numbers of 128K (0.31 Hz/point) and 64K (0.61 Hz/point) were used for the 8 and 32 Hz linewidth data series, respectively, while a Fourier number of 16K (2.4 Hz/point) was used for both the 240 and 832 Hz linewidth data series. Time-domain S/N ratios were measured by comparing the signal amplitude in the combined data to the RMS noise in the last 5% of the data. Frequency-domain S/N ratios were estimated using Varian software that calculates the S/N ratio as the signal height divided by 2 × RMS noise level. All calculations and data processing were performed on SUN workstations using either automated C-shell scripts or Varian MAGICAL macros.

Discrete Fourier transformation procedures. Each combined data set for the four different HOD samples was analyzed using Varian software (VNMR software version 3.1) to obtain frequency and signal amplitude estimates of the HOD resonance according to the following four DFT procedures: (i) DFT-I, unweighted (i.e., no apodization) with a software-driven automatic zero-order phase correction (i.e., without linewidth and signal phase prior information, but see below); (ii) DFT-II, unweighted but using a software-determined zero-order phase correction obtained from the highest S/N data set (i.e., without linewidth prior information, but with signal phase prior information); (iii) DFT-III, weighted with a matched exponential filter and using a software-driven automatic zero-order phase correction (i.e., with linewidth prior information, but without signal phase prior information); and (iv) DFT-IV, weighted with a matched exponential filter and using a software-determined zero-order phase correction obtained from the highest S/N data set (i.e., with both linewidth and signal phase prior information). The DFT analysis procedures are summarized in Table 1.

TABLE I

Summary of DFT and Bayesian Analysis Procedures

Method	Linewidth prior information	Signal phase prior information	Labels (plots, curve)
DFT-I[a]	No	No	A, a
DFT-II	No	Yes	C, c
DFT-III	Yes	No	E, e
DFT-IV	Yes	Yes	G, g
Bayesian-I	No	No	B, b
Bayesian-II	No	Yes	D, d
Bayesian-III	Yes	No	F, f
Bayesian-IV	Yes	Yes	H, h

[a] A matched filter was applied before using automatic zero-order phasing routines, but then removed to obtain frequency and signal amplitude parameter estimates.

Off-resonance frequency estimates at a given σ for the combined data sets were obtained by picking the highest point in the frequency-domain NMR absorption spectrum. The integrated intensity or signal amplitude estimates were determined by taking the Gaussian quadrature integral or straight sum over the region of the absorption spectrum containing the peak. Integration limits for this region were adjusted for each HOD sample to encompass 90% of the total integrated intensity for a Lorentzian lineshape [i.e., $\pm 6.3 \times$ (full width at half-height of the highest S/N data)]. This represents a hidden use of prior information. When matched weighting filters and zero-order phase corrections were used as prior information in the DFT-II, DFT-III, and DFT-IV analysis procedures, they were determined from the highest S/N data set in the series using analytical procedure DFT-I.

A number of ancillary measures were taken to help improve the robustness of the software-driven DFT procedures. To improve the automatic spectral phasing at low S/N levels, a matched filter (prior information) was applied before using automatic phasing routines in procedure DFT-I, but then removed for purposes of estimating the NMR parameters. Additional corrections for errors introduced by the automatic phasing routine were made for all the DFT procedures by applying a 180° phase correction when a negative signal amplitude was found within $\pm$ linewidth hertz of the known frequency. Considerable improvement in the signal amplitude estimates at low S/N levels was obtained by setting the integration limits on the basis of the known linewidth (including contributions resulting from line-broadening apodization), rather than the linewidth determined for the estimated resonance by the DFT procedures. These modifications are equivalent to using prior information and introduce a form of subjectiveness to the analysis procedure that mimics that which would accompany manual processing techniques. Without such modifications, *much greater* excursions in the precision of the parameter estimates were observed at high σ values using the DFT procedures.

Bayesian analysis procedures. For the Bayesian analysis, the series of combined data sets for the four HOD samples were analyzed according to the following four

procedures, which correspond to similar procedures used in the DFT analysis: (i) Bayesian-I, independent of the signal decay rate constant and phase; (ii) Bayesian-II, independent of the signal decay rate constant, but given the signal phase; (iii) Bayesian-III, given the signal decay rate constant, but independent of the signal phase; and (iv) Bayesian-IV, given both the signal decay rate constant and the signal phase. The Bayesian analysis procedures are summarized in Table 1.

The posterior probability distributions for the frequencies were calculated independent of the signal amplitudes, while the posterior probability distributions for the signal amplitudes were calculated independent of the frequencies. Decay rate constant and signal phase prior information was incorporated in the Bayesian-II, Bayesian-III, and Bayesian-IV analysis procedures in a manner similar to that reported previously (*17*). Inclusion of prior information directly into the Bayesian estimation procedures eliminates the need to treat the decay rate constant as a nuisance parameter. Such modifications greatly reduce the computational time by eliminating the need to marginalize the decay rate constant using the rules of probability theory (*22*). For all amplitude estimates, the frequency was marginalized. A detailed description and theoretical basis for these procedures can be found in the literature (*17–22*). In addition to estimating the parameters, probability theory carries with it an estimate of the uncertainty of the parameter estimate (*26*), which can be determined from the width of the posterior probability distributions.

RESULTS AND DISCUSSION

This study compares the ability of DFT and Bayesian probability theory to estimate the frequency and the signal amplitude of a single, exponentially decaying sinusoid from NMR FID data. Comparisons are made as a function of the S/N level of the FID data. Parameter estimates obtained for very high S/N FID data are taken as the true parameter values (frequency, -1000 Hz; normalized signal amplitude, 100). The deviation or scatter from these values reflects uncertainty in the estimates introduced by the addition of noise.

Linewidth dependence. Figure 1 shows representative ^{1}H NMR spectra and their frequency-domain S/N values for all four HOD samples, which were obtained using the DFT-I analytical procedure at two time-domain S/N levels ($\sigma = 1$ and $\sigma = 10$). It is clear from Fig. 1 that identical time-domain S/N ratios at a given σ do not translate into identical frequency-domain S/N ratios. The height of the HOD resonance obtained from the absorption spectrum scales inversely with the resonance linewidth, thus the frequency-domain S/N ratio will decrease as the decay rate constant increases for constant time-domain S/N levels. Such a dependence upon linewidth has profound effects upon the performance of both the DFT and the Bayesian methods.

The average frequency and signal amplitude estimates ($\pm$one standard deviation) obtained using the DFT-IV procedure for the 8 and 832 Hz linewidth HOD samples are plotted as a function of σ in Fig. 2. The plots reveal that both linewidths display nearly identical functional behavior as the S/N level decreases. Estimation failure occurs at the point where the standard deviation in the frequency estimates dramatically increase. Failures are observed at $\sigma \cong 350$ and $\sigma \cong 50$ for the 8 Hz (Fig. 2A) and 832 Hz (Fig. 2B) linewidth HOD samples, respectively, and represent the S/N levels beyond which the technique becomes essentially useless for estimating frequencies. These σ

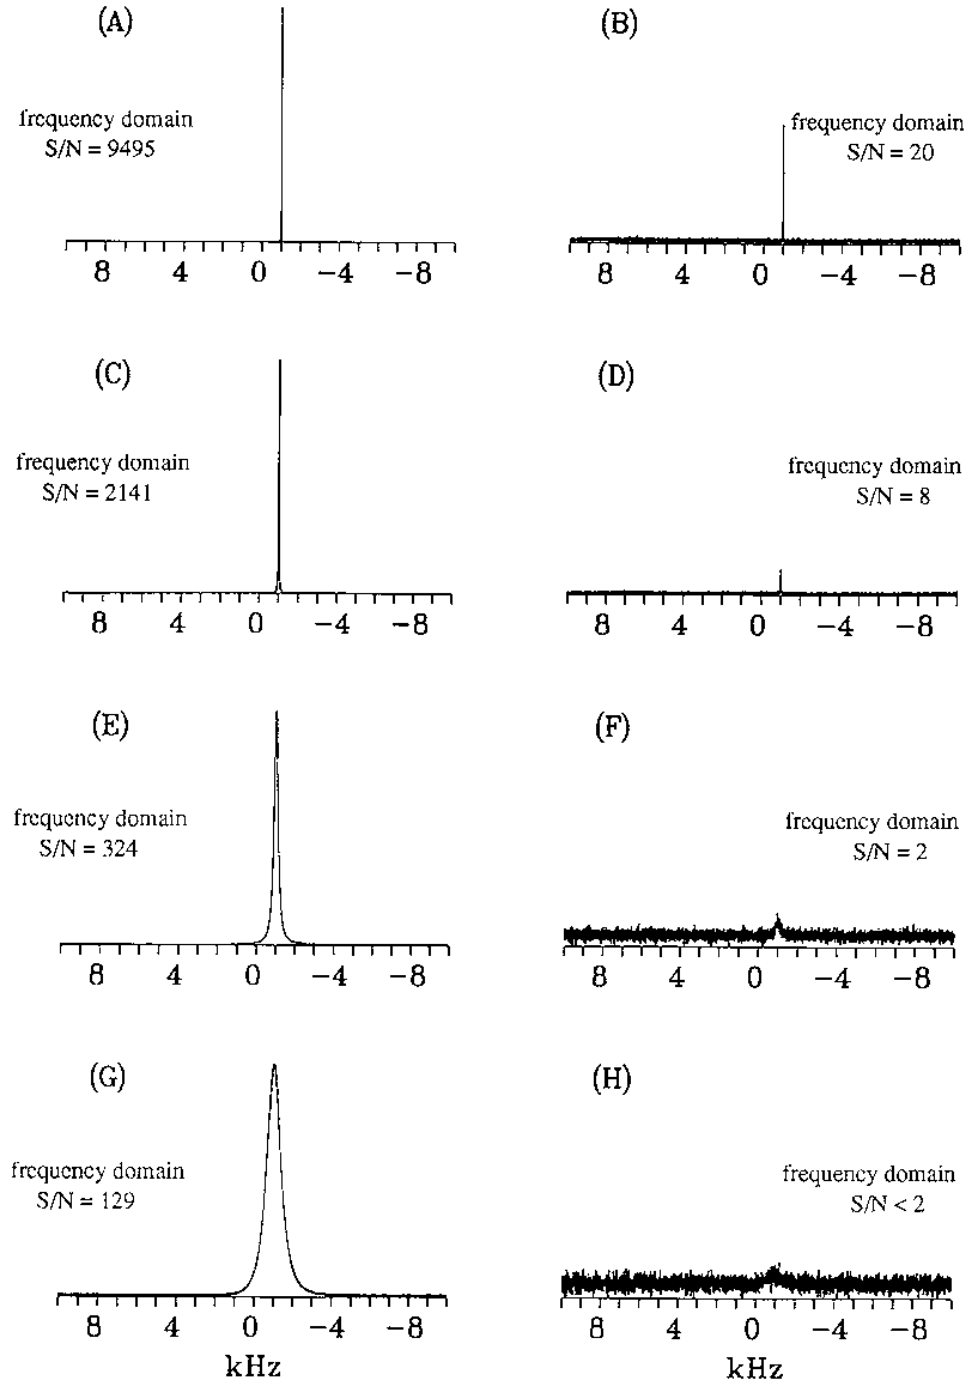

FIG. 1. Representative ^{1}H NMR absorption spectra for the 8 Hz (A and B), the 32 Hz (C and D), the 240 Hz (E and F), and the 832 Hz (G and H) HOD samples for noise standard deviations (σ) of 1 and 10. The spectra were generated using the DFT-I analysis procedure (i.e., no exponential apodization with an automatic zero-order phase correction). While the integrated area in each spectrum is identical, vertical scales have been adjusted for each plot so all peak heights are equal (A, C, E, and G) or are visible (B, D, F, and H).

values correspond to time-domain S/N ratios of approximately 0.3 and 2.0, respectively. Above these failure points, peak-picking routines report the highest random noise peak in the absorption spectrum as the frequency. Unlike the frequency-estimation behavior, the DFT-IV analysis yields average signal amplitude curves in Fig.

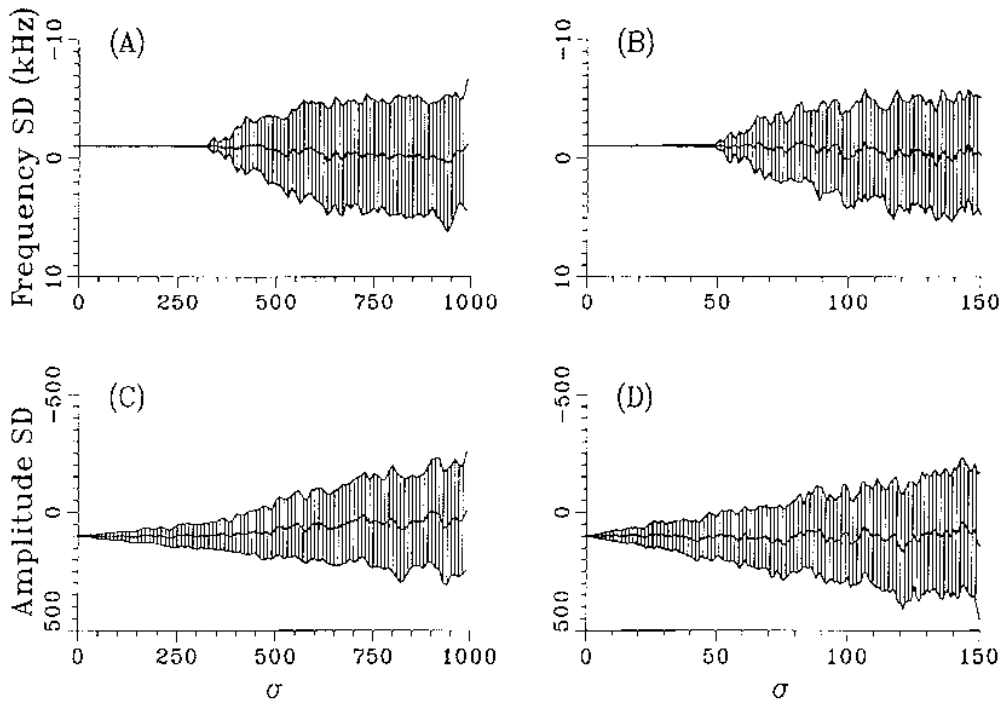

FIG. 2. The average frequency and the signal amplitude estimates (±one standard deviation) for the 8 Hz (A and C) and 832 Hz (B and D) HOD samples plotted as a function of the noise standard deviation (σ). Average parameter values and standard deviations were computed using the DFT-IV procedure (i.e., with prior linewidth and signal phase information) at each σ using 50 individual combined data sets (signal plus noise). Estimated frequencies and amplitudes are referenced or normalized to the highest S/N data set (frequency, 1000 Hz; normalized signal amplitude, 100).

2 that show no dramatic failure point for either linewidth. Errors in the amplitude estimates increase linearly with increases in σ. At high σ values, the average signal amplitude estimates slowly approach zero. The tendency toward zero at poor S/N levels is a consequence of averaging amplitude estimates resulting from individual noise peaks which have been selected as the resonance frequency.

The data displayed in Fig. 2, along with similar plots for the 32 and 240 Hz linewidth HOD samples (data not shown), demonstrate that the uncertainty in the parameter estimates and the failure point for the DFT-IV frequency estimate scales with the NMR resonance linewidth. Similar behavior was verified for the DFT-I, DFT-II, DFT-III, and all four Bayesian procedures. For a given S/N level, an increase in the linewidth of the NMR signal increases the uncertainty in the NMR parameter estimates. For similar decreases in the S/N level, however, the percentage of change in the precision for a given analysis procedure is independent of the linewidth. Comparison of all linewidth data indicates that the frequency-estimation failure point scales roughly inversely with the square root of the signal decay rate. In summary, although the linewidth of the NMR signal affects the precision attained at a fixed time-domain S/N level for a given DFT or Bayesian analysis procedure, differences in linewidth merely scale the results for a given method and analytical procedure. Therefore, further discussions focus upon a single linewidth with the understanding that results for other linewidths have the same functional form and are scaled accordingly.

Frequency estimation. Figure 3 summarizes the individual DFT and the Bayesian frequency estimates for a 32 Hz linewidth HOD sample plotted as a function of σ.

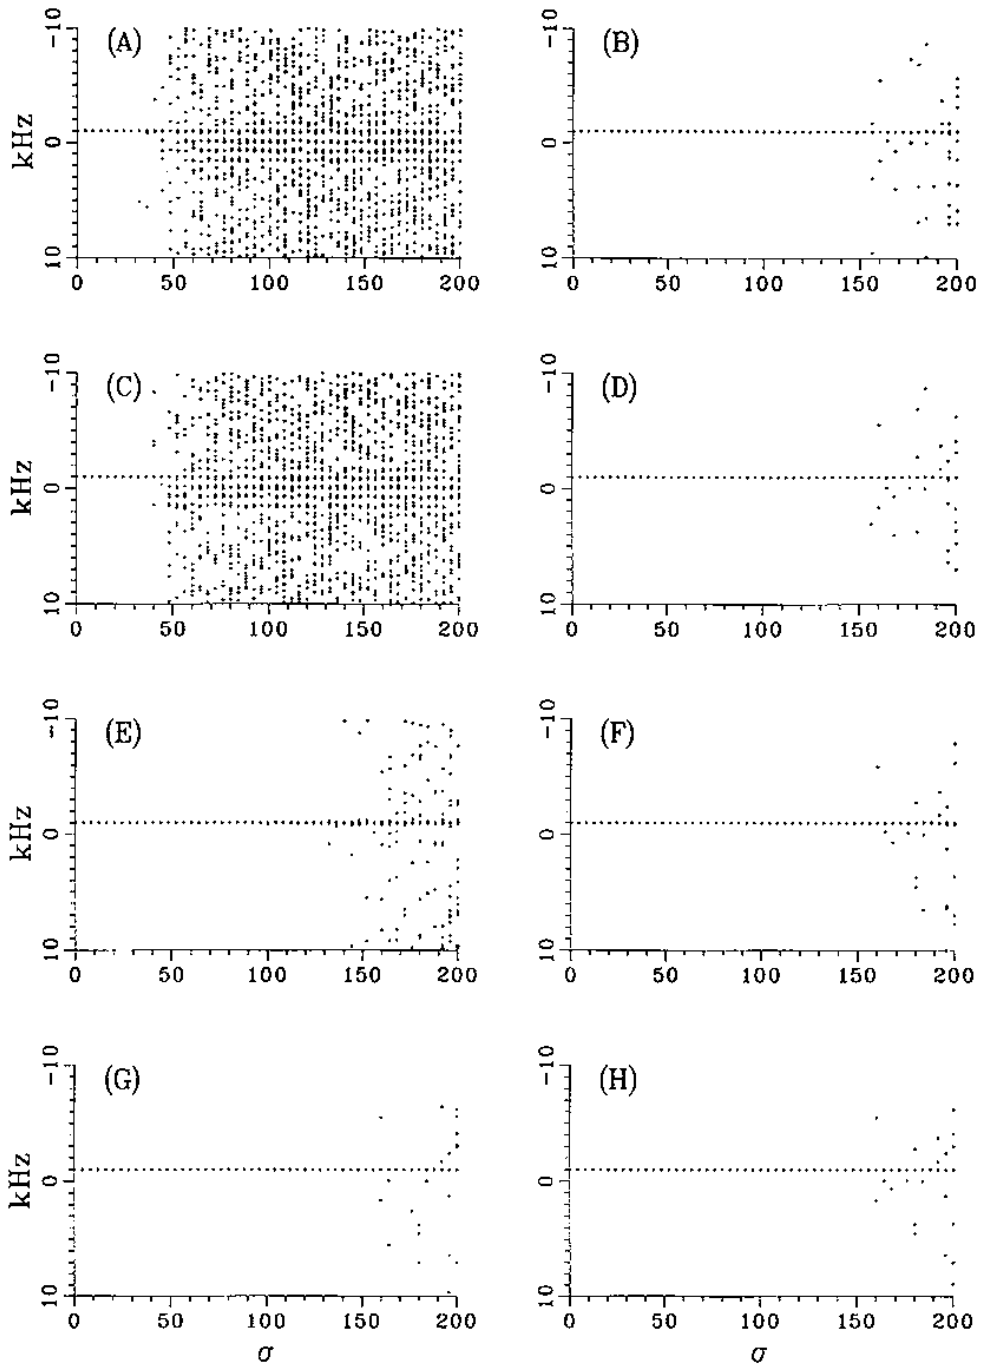

FIG. 3. Individual frequency estimates of 32 Hz linewidth HOD sample plotted as a function of the noise standard deviation (σ). Results are shown for (A) the DFT-I, (B) the Bayesian-I, (C) the DFT-II, (D) the Bayesian-II, (E) the DFT-III, (F) the Bayesian-III, (G) the DFT-IV, and (H) the Bayesian-IV analyses procedures. (See text and Table I for description of the procedures). Fifty independent frequency estimates were made at each σ value for a total of 2500 individual estimates in each panel.

Direct comparisons and behavioral trends for all four DFT and Bayesian analysis procedures can be made as σ increases by examining the amount of scatter in the data. All analysis procedures for both methods give frequency estimates with minimal scatter for the HOD resonance at high S/N levels. At lower S/N levels, the difference

in the amount of scatter observed for the DFT and Bayesian methods is striking and depends on the prior information supplied during the analysis.

The scatter in the frequency estimates obtained using the DFT-I and the DFT-II analysis procedures (Figs. 3A and 3C, respectively) at σ values <30 is negligible on the scale of these plots, but increases dramatically at higher σ values. Use of signal phase prior information has little effect on the amount of scatter in the DFT frequency estimates, reflecting good performance for the modified automatic phasing routines (i.e., phasing in the presence of a matched exponential-weighting filter). Significant improvements are achieved in the DFT-III and DFT-IV procedures (Figs. 3E and 3G) when prior knowledge of the linewidth permits the appropriate matched weighting filter to be applied to the data during the parameter-estimation procedures. As expected, the least amount of scatter in DFT frequency estimates is obtained using the DFT-IV procedure when prior information on both the linewidth and the signal phase is supplied (Fig. 3G) in the analysis.

In contrast to the DFT results, the Bayesian frequency-estimation results show little dependence upon the direct incorporation (as opposed to treatment as nuisance parameters) of signal decay rate constant and signal phase prior information into the analysis procedure. Such findings are not unexpected, since use of prior information is an *inherent* feature of the Bayesian analysis, even where it is not directly supplied. As a result, all four of the Bayesian procedures (Figs. 3B, 3D, 3F, and 3H) give nearly identical results. Scatter in the Bayesian results is significantly reduced from the scatter observed for the DFT-I and DFT-II procedures. The Bayesian frequency estimates are also superior to the DFT-III frequency estimates. *In accordance with theory, however, when both a matched weighting filter and the known signal phase are used, the frequency estimates and scatter obtained using the Bayesian analysis (Fig. 3H) and the DFT method (Fig. 3G) are essentially identical.*

Although considerable effort was taken to improve the performance of the software-driven DFT procedures (*vida supra*), the methods fail to estimate the HOD frequency when the highest peak in the absorption spectrum becomes a uniformly distributed noise peak. At this point the frequency estimates begin to randomly deviate from the known value of -1000 Hz. The σ value at the failure point can be approximated for the 32 Hz linewidth HOD sample from the scatter plots in Fig. 3. The DFT-I and DFT-II procedures begin to fail at σ values between 40 and 50, while the DFT-III procedure fails at $\sigma \cong 130$. The failure point in the DFT analysis is greatly affected by the use of decay rate constant and signal phase prior information. Apodization of the FID NMR data using a matched weighting filter extends the failure point to higher σ by effectively increasing the frequency-domain S/N of the absorption spectrum and, thus, improves the utility of simple peak-picking procedures. By contrast, application of prior signal phase information, compared to the use of modified automatic phasing routines, has little effect on the failure point of the DFT analysis. Since the degree of scatter in the analysis depends upon the decay rate of the NMR signal (see Fig. 1), the failure point expected for the DFT-IV analysis of the 32 Hz linewidth HOD sample (Fig. 3G) should fall between the values determined for the 8 and 832 Hz linewidth HOD samples ($\sigma = 50$ and $\sigma = 350$, respectively). This is observed. The DFT-IV frequency estimates for the 32 Hz linewidth HOD sample start to fail at $\sigma \cong 150$, consistent with our earlier observation that parameter estimates for fast-decaying sig-

nals are less precise than those for slowly decaying signals for a given time-domain S/N level.

Unlike DFT frequency estimates, which depend upon the frequency-domain S/N level, Bayesian procedures depend upon the time-domain S/N level. Bayesian frequency estimates are not subject to errors introduced by the peak-picking routines applied in the frequency domain. Consequently, different failure points are observed for the Bayesian methods and for the DFT methods. Bayesian analysis begins to fail when the spectral width of the posterior probability distribution for the frequency becomes comparable to the spectral width of the NMR absorption spectrum. Figure 3 demonstrates for all the Bayesian analytical procedures that this failure occurs at S/N levels much lower ($\sigma \cong 160$) than those in the DFT-I, DFT-II, and DFT-III procedures. In fact, at σ values where the DFT analysis procedures fail completely, the time-domain S/N ratio is still sufficient to provide reasonable frequency estimates using Bayesian methodology. In the absence of decay-rate-constant prior information, the Bayesian method is more computationally intensive than the DFT method, requiring on the order of 20 times more processing time. When this information is supplied, however, the amount of computational time required for the two procedures is approximately the same. Such reductions in processing time are achieved by eliminating the need to marginalize nuisance parameters and have important consequences in practical applications of the Bayesian method, where one can often estimate the NMR resonance linewidth.

Signal amplitude estimations. Figure 4 shows the individual DFT and Bayesian signal amplitude estimates for the same 32 Hz-linewidth HOD sample that was examined in Fig. 3. At high S/N levels ($\sigma < 30$), moderate scatter in the signal amplitude estimates is observed for both methods. Considerably more scatter and large differences between the two methods are noted at low S/N levels. The most dramatic difference between the DFT and Bayesian procedures is seen as groups of estimates centered around zero amplitude in the DFT-I and DFT-II signal amplitude plots (Figs. 4A and 4C, respectively). In the absence of a matched filter, these methods begin to fail at estimating the signal amplitude at a σ value between 40 and 50. This point corresponds well with the failure point observed for the DFT-I and DFT-II frequency estimates in Fig. 3. Use of signal phase information (Fig. 4C) gives only slight improvements in the DFT amplitude estimates, again indicating that the automatic phase calculations are working well in the presence of a matched-exponential filter. More dramatic improvements are observed in the DFT signal amplitude estimates when a matched filter is applied to the data (Figs. 4E and 4G) during the parameter-estimation procedures. The bands of data points at zero amplitude essentially disappear. As with the frequency estimates, the smallest amount of scatter for the DFT amplitude estimates is observed when prior information on both the decay rate constant and the signal phase is used in the analysis (Fig. 4G).

The Bayesian signal amplitude estimates shown in Fig. 4 exhibit less scatter than the DFT amplitude estimates. Like the Bayesian frequency measurements, direct incorporation of the signal decay rate constant and signal phase prior information into the analysis procedure (e.g., procedures II, III and IV) has little effect on the scatter in the Bayesian amplitude estimates, though slight improvements at low S/N levels are observed when prior knowledge of the signal decay rate is provided. This contrasts

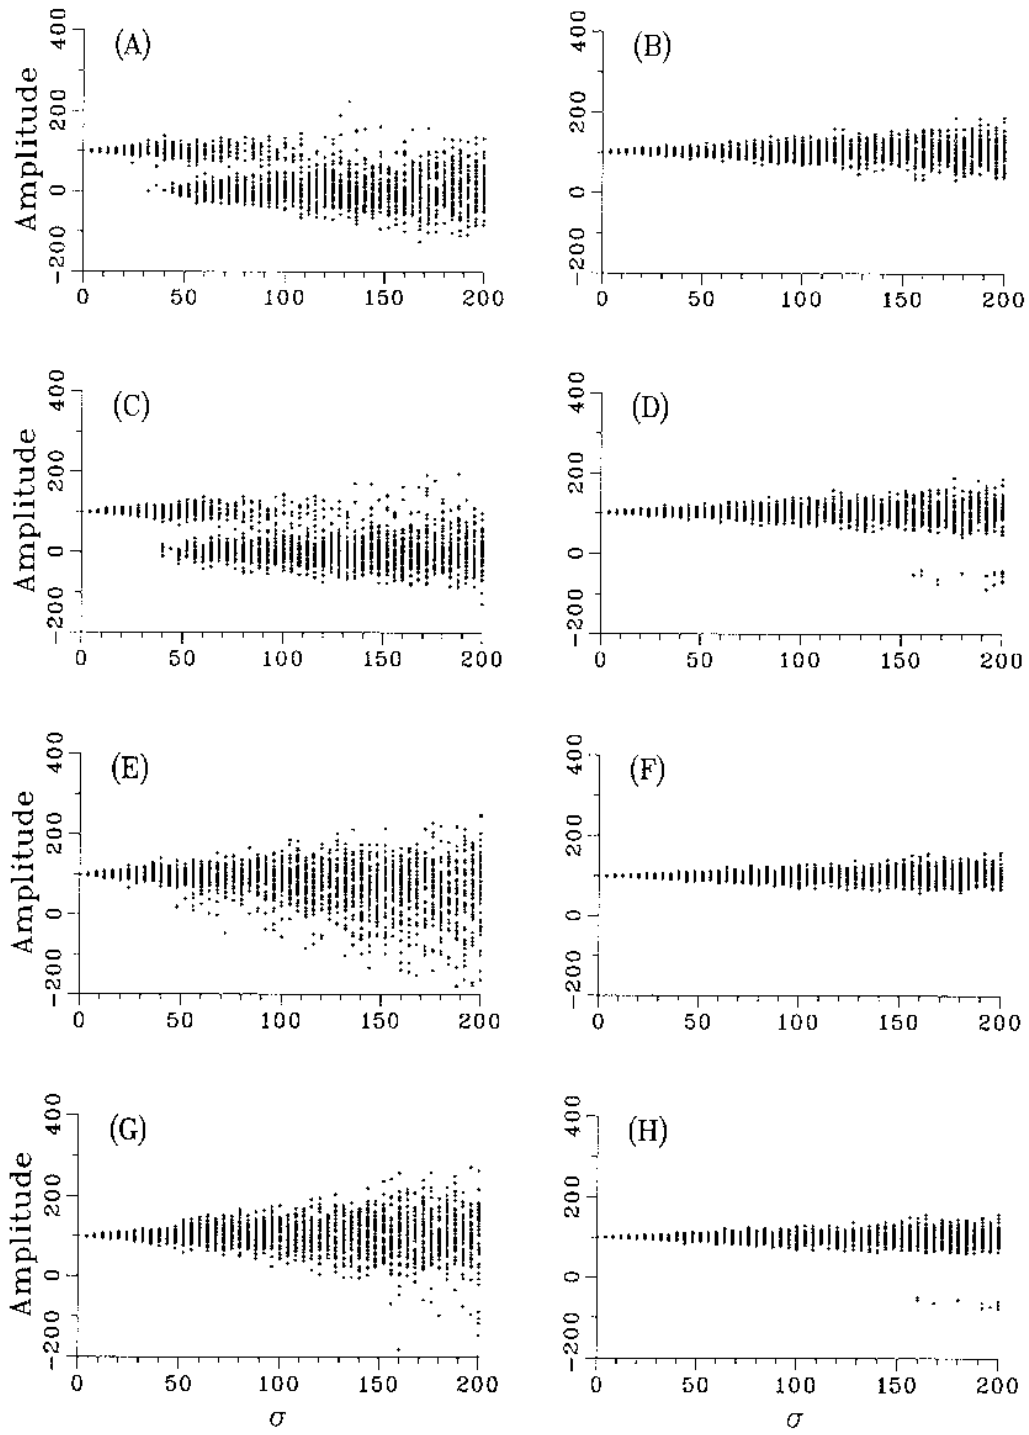

FIG. 4. Individual signal amplitude estimates of 32 Hz linewidth HOD sample plotted as a function of the noise standard deviation (σ). Results are shown for (A) the DFT-I, (B) the Bayesian-I, (C) the DFT-II, (D) the Bayesian-II, (E) the DFT-III, (F) the Bayesian-III, (G) the DFT-IV, and (H) the Bayesian-IV analysis procedures. (See text and Table 1 for description of the procedures). Fifty independent amplitude estimates were made at each σ value for a total of 2500 individual estimates in each panel.

with the DFT amplitude estimates, which show a strong dependence upon the prior information. In addition, the groups of amplitude estimates centered around zero (Figs. 4A and 4C) are not present (Figs. 4B and 4F) or are present only at $\sigma > 150$ (Figs. 4D and 4H). Comparison of the Bayesian and DFT results in Fig. 4 shows that

the Bayesian procedures yield significantly better estimates of the signal amplitude at low S/N levels than any of the DFT analysis procedures. In many cases, Bayesian analysis provides reasonable amplitude estimates under conditions where the DFT method completely fails, making it the preferred choice for estimating the signal amplitude.

The failure of the DFT-I and DFT-II procedures to estimate the signal amplitude at $\sigma > 40$ is due to the inability of the peak-picking routines to distinguish between peaks and noise spikes. Amplitudes for a given signal can only be estimated and are meaningful only when the peak-picking routine returns a reasonable frequency estimate. If a noise spike is identified as the frequency, for instance, a meaningless random value close to zero is estimated for the signal amplitude. Such effects force the average amplitude estimate (obtained from a number of individual estimates) toward zero, as shown in Figs. 2C and 2D, and accounts for the group of signal amplitude estimates located at zero in Figs. 4A and 4C. Changes in the DFT processing procedures which improve the DFT frequency estimates will also lead to improvements in the signal amplitude estimates. On the other hand, Bayesian methodology is not limited by the robustness of the frequency-domain peak-picking routine. Direct use of prior information of the decay rate constant and the signal phase gives essentially the same results as those when this information is not provided. As for the Bayesian frequency-estimation procedures, a 20-fold reduction in processing time is achieved by using the decay rate constant as prior information. Smaller reductions would be expected when frequency prior information is used during the amplitude-estimation Bayesian analysis.

Parameter-estimate uncertainties. The frequency and signal amplitude standard deviations computed for all analytical procedures at each σ for the 32 Hz linewidth HOD sample are plotted as a function of σ in Fig. 5. The parameter standard deviations provide an alternate representation of the data shown in Figs. 3 and 4 and reflect inherent differences between the DFT and the Bayesian methods. This figure permits an assessment of the precision of the NMR parameter estimates for each of the various analytical procedures as a function of the S/N level. In general, Bayesian methods outperform the DFT methods at all S/N levels and continue to provide parameter estimates well after DFT methods fail due to poor S/N levels.

Figure 5A shows the frequency standard deviation curves for all the DFT and Bayesian analytical procedures. With the exception of the DFT-IV analysis (curve g), the DFT procedures (curves a, c, and e) give uncertainties in the frequency estimates larger than those of any of the corresponding Bayesian procedures (curves b, d, f, and h) at all S/N levels. Superposition of curves d, g, and h demonstrates the similarities in the performance of the DFT-IV and the Bayesian-II and Bayesian-IV procedures (curves d and h, respectively). Frequency-estimation failure points for each procedure shown in Fig. 5A are clearly marked by a large increase in the frequency standard deviation (from less than 25 Hz to approximately 5000 Hz). Larger values for the frequency standard deviation past these points are not plotted. Up to the failure point, a linear increase in the standard deviation of the parameter estimates is observed for a linear increase in σ. Consistent with observations made for the individual frequency estimates in the scatter plots, the failure points for the DFT-I and DFT-II procedures (curves a and c) occur at $\sigma \cong 30$ and for the DFT-III analysis (curve e) at $\sigma \cong 130$. The DFT-IV and all of the Bayesian procedures demonstrate a performance better

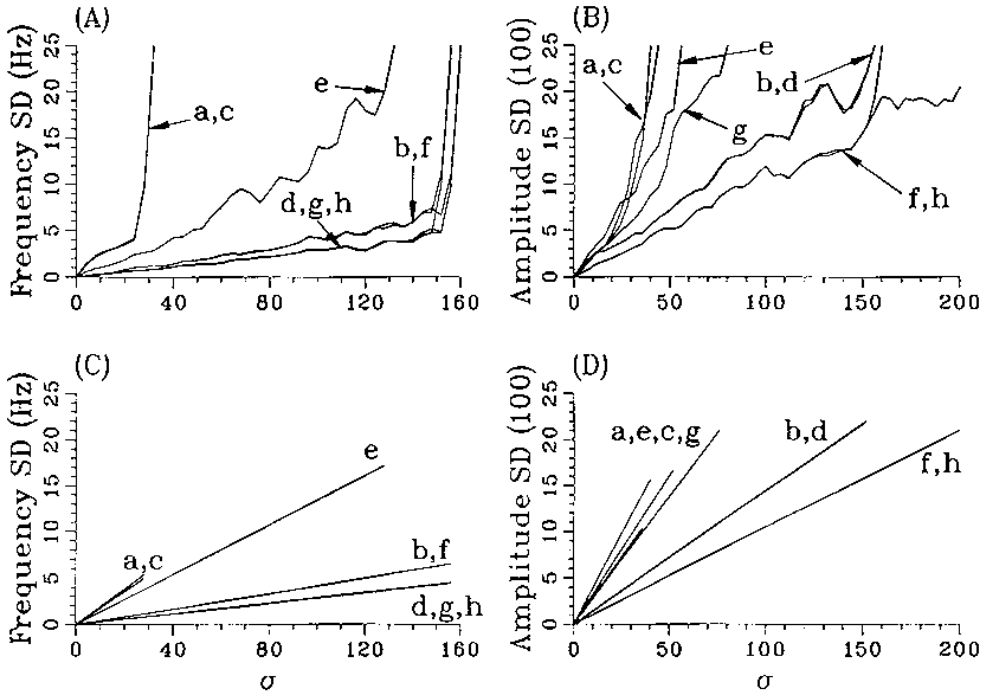

FIG. 5. The standard deviations for the (A) frequency and (B) signal amplitude estimates for the 32 Hz HOD sample plotted as a function of the noise standard deviation (σ). NMR parameter standard deviations were computed at each noise standard deviation from the 50 independent measurements shown in the scatter plots in Fig. 3 and Fig. 4, respectively. The curves are designated according to the analytical procedure that was used to obtain the parameter estimate as shown in Table 1: a, DFT-I; b, Bayesian-I; c, DFT-II; d, Bayesian-II; e, DFT-III; f, Bayesian-III; g, DFT-IV; and h, Bayesian-IV. Lines have been fit to the initial portions of the curves in (A) and (B) to yield (C) and (D), respectively.

than those of the other DFT procedures at low S/N levels by extending the failure point for the frequency estimates to $\cong$150.

Likewise, the curves in Fig. 5B show that the Bayesian method provides precision for estimating signal amplitudes better than that of the DFT method. The most dramatic improvements are observed for σ values >30. Failure points for the amplitude-estimate curves are not as clearly defined as those in the corresponding frequency-estimate curves (Fig. 5A), but still occur where a substantial increase in the amplitude standard deviation is observed. Similar to the frequency estimates, all procedures show a linear increase in the amplitude standard deviation with increasing σ. It is clear that the Bayesian procedures provide reasonable amplitude estimates at S/N levels where the DFT procedures have failed. In fact, all of the DFT procedures fail at σ values smaller than those observed for the Bayesian procedures. The DFT-I and DFT-II procedures (curves a and c) fail at $\sigma \cong 30$, the DFT-III procedure (curve e) fails at $\sigma \cong 50$, and the DFT-IV procedure (curve g) fails at $\sigma \cong 80$. The Bayesian-I, Bayesian-II, and Bayesian-III amplitude estimates (curves b, d, and f) fail at about the same σ value for all procedures ($\sigma \cong 150$). The Bayesian-IV analysis is slightly better, failing at $\sigma > 170$.

In practice, meaningful DFT amplitude estimates can be obtained only at S/N levels greater than those set by the failure points for the frequency estimates (i.e., one cannot estimate an amplitude for a signal if its frequency is undetermined). The σ

value at the frequency failure point represents an upper limit beyond which the DFT *amplitude* estimates are meaningless. In contrast, reasonable *frequency* estimates using the DFT procedures can be obtained at σ values where amplitude estimates have clearly failed. For example, both DFT-I and DFT-II amplitude estimation begins to fail at approximately the same σ value where DFT-I and DFT-II frequency estimation starts to fail ($\sigma \cong 30$). Amplitude estimates for the DFT-III and DFT-IV procedures, on the other hand, fail at much lower σ values than the frequency estimate for these procedures. Similarly, extraction of meaningful Bayesian amplitude estimates are limited by S/N levels. When phase prior information is available, amplitude-estimation failure points always occur at σ values similar to those for the Bayesian frequency estimates. When decay rate constant prior information is used, however, no sharp failure point for the Bayesian amplitude estimate is apparent. Instead a smooth increase in the uncertainty is observed even when the frequency estimates begin to fail.

A more quantitative comparison of the precision for the NMR parameter estimates obtained from the DFT and Bayesian methods can be made by fitting lines (linear least squares) to the initial portions of the curves in Figs. 5A and 5B. Fits are shown in Figs. 5C and 5D for the frequency and signal amplitude estimates, respectively. Such a fitting procedure is reasonable, given the linear relationship that exists (up to the failure point) between σ and the uncertainty in the parameter estimates. The slopes are summarized in Table 2 and establish the relative order of performance for the DFT and Bayesian analytical procedures. Smaller slopes correspond to better precision in the parameter estimate. The order of performance established from Table 2 for the DFT and Bayesian frequency estimates is DFT-IV $\cong$ Bayesian-II $\cong$ Bayesian-IV $>$ Bayesian-I $\cong$ Bayesian-III $>$ DFT-III $\gg$ DFT-I $\cong$ DFT-II. A similar ranking of the different analytical procedures can be established for the signal amplitude estimates from the slopes in Fig. 5D. The order of performance for the signal amplitude estimates is Bayesian-III $\cong$ Bayesian-IV $>$ Bayesian-I $\cong$ Bayesian-II $>$ DFT-IV $>$ DFT-II $>$ DFT-III $>$ DFT-I. Examination of similar curves for the 8, 240, and 832 Hz linewidth HOD samples (data not shown) shows that the order of performance for the different

TABLE 2

Relative Order of Performance for DFT and Bayesian Analysis Procedures Reported as Slopes for Lines in Fig. 5

Analysis method	Frequency (SD slope)	Signal amplitude (SD slope)
DFT-I	0.19	0.39
DFT-II	0.18	0.29
DFT-III	0.13	0.32
DFT-IV	0.029	0.28
Bayesian-I	0.041	0.14
Bayesian-II	0.029	0.14
Bayesian-III	0.041	0.11
Bayesian-IV	0.028	0.10

analytical procedures is independent of the NMR signal decay rate constant. An increase in the decay rate constant results in corresponding increases in the slopes for the frequency and signal amplitude lines shown in Figs. 5C and 5D, while a decrease in the decay rate constant reduces the value of the slope. Changes in the parameter-estimation failure point for each analytical procedure are also scaled by the decay rate constant of NMR signal. In no case, however, does the performance order for the different analytical procedures change. As previously discussed, the direct use of prior information also affects the performance of the DFT analysis and to a lesser extent the performance of the Bayesian analysis. The slopes in Table 2 show that with analysis methods I, II, and III, Bayesian methods provide a factor of 3 to 7 improvement in the precision of the frequency estimates at S/N levels where the DFT procedures are successful. Even greater improvements are obtained using the Bayesian method at S/N levels where the DFT procedures fail. Table 2 also shows that at all S/N levels, Bayesian analysis is better at estimating the signal amplitude than the DFT procedures by a factor of 2 to 3.

A major disadvantage with DFT analysis of NMR FID data is that no information about the uncertainty in the estimate can be obtained for a single measurement. In contrast, Bayesian analysis provides a direct measure of the uncertainty in the parameter estimate from the width of the probability distribution (*26*). These widths are plotted for the Bayesian-I frequency and amplitude estimates of the 240 Hz linewidth HOD sample in Figs. 6A and 6B, respectively. To determine if the widths of the posterior are a good measure of the uncertainty in the estimate, the difference between the known parameter value and the estimate was computed. The results showed the true parameter value to be distributed within one width of the posterior for 68% of the time, within two widths of the posterior for 93% of the time, and within three widths of the posterior for 99% of the time (i.e., a Gaussian distribution). Thus, the widths of the Bayesian probability distribution are good indications of the uncertainty in the parameter estimates.

SUMMARY AND CONCLUSIONS

Our results demonstrate that Bayesian probability theory offers distinct advantages over the DFT for estimating frequencies and amplitudes of NMR signals from time-

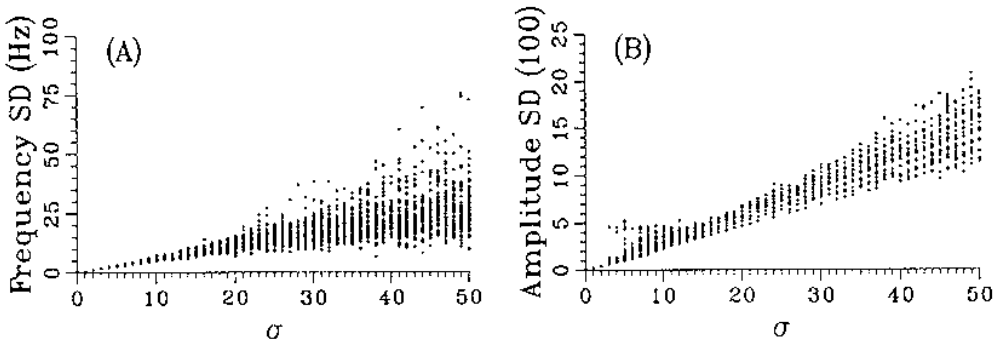

FIG. 6. Plot of the width of the posterior probability distribution *versus* the noise standard deviation (σ) for (A) the frequency and (B) the signal amplitude estimates for the 240 Hz HOD sample. Fifty individual points at each σ value are shown.

domain FID data. These advantages have been demonstrated for the single-frequency case and are equally applicable to multiple, well-separated frequencies. Our conclusions are based on measured uncertainties determined from repetitive estimates of the NMR parameters obtained using a variety of analytical procedures.

Evaluation of Bayesian and DFT methods, with and without prior information, has indicated that (a) the DFT analysis is very sensitive to whether prior information is included while the Bayesian analysis is much less sensitive; (b) Bayesian analysis provides more precise estimates of the frequency than the DFT method, except when prior information on both the decay rate constant and phase of the NMR signal is supplied during data analysis, in which case the two methods yield essentially identical *frequency* estimates with similar precision; (c) Bayesian analysis gives more precise estimates of the signal amplitude than the DFT method; (d) the order of performance for the DFT and Bayesian analytical procedures (shown in Table 2) is independent of the decay rate of the NMR signal; and (e) the Bayesian method provides a measure of the precision for a single estimate of the NMR parameter for which no comparable measure can be obtained using DFT methodology.

As with other time-domain analysis techniques, Bayesian analysis does not rely on secondary parameter extraction procedures which examine the frequency-domain NMR spectrum. Instead Bayesian analysis estimates the parameters of interest by direct examination of the time-domain data. Under practical conditions where the signal decay rate constant and phase are unknown (or a distribution of these is present), more precise parameter estimates can be expected using Bayesian analysis than the DFT analysis. We anticipate that Bayesian procedures will play significant roles as alternative, time-domain NMR data analysis techniques. Our future investigations will include overlapping multiple frequencies, truncated FIDs, and multidimensional NMR data.

ACKNOWLEDGMENTS

We thank Joel R. Garbow for helpful discussions. This work was supported by the Monsanto Co. and, in part, by NIH Grant GM-30331, J. J. H. Ackerman principal investigator.

REFERENCES

1. J. Lowe and R. E. Norberg, *Phys. Rev.* **107,** 46 (1957).
2. J. W. Cooley and J. W., *Math. Comput.* **19,** 297 (1965).
3. R. R. Ernst and W. A. Anderson, *Rev. Sci. Instrum.* **37,** 93 (1966).
4. J. C. Lindon and A. G. Ferrige, *Prog. NMR Spectrosc.* **14,** 27 (1980).
5. D. Shaw "Fourier Transform NMR Spectroscopy," 2nd ed., Elsevier, Amsterdam, 1984.
6. R. R. Ernst, G. Bodenhausen, and A. Wokaun, "Principles of Nuclear Magnetic Resonance in One and Two Dimensions," Oxford Univ. Press, Oxford, 1987.
7. G. Otting, H. Wider, G. Wagner, and K. Wüthrich, *J. Magn. Reson.* **66,** 187 (1986).
8. D. S. Stephenson, *Prog. NMR Spectrosc.* **20,** 515 (1988).
9. S. Sibisi, *Nature* **301,** 134 (1983).
10. S. Sibisi, J. Skilling, R. G. Brereton, E. D. Laue, and J. Staunton, *Nature* **311,** 446 (1984).
11. E. D. Laue, J. Skilling, J. Staunton, S. Sibisi, and R. G. Brereton.73 *J. Magn. Reson.* **62,** 437 (1985).
12. J. F. Martin, *J. Magn. Reson.* **65,** 291 (1985).
13. F. Ni, G. C. Levy, and H. A. Scheraga, *J. Magn. Reson.* **66,** 385 (1986).
14. A. R. Mazzeo and G. C. Levy, *Magn. Reson. Med.* **17,** 483 (1991).

15. J. A. JONES AND P. J. HORE, *J. Magn. Reson.* **92,** 276 (1991).
16. J. A. JONES AND P. J. HORE, *J. Magn. Reson.* **92,** 363 (1991).
17. G. L. BRETTHORST, C.-C. HUNG, D. A. D'AVIGNON, AND J. J. H. ACKERMAN, *J. Magn. Reson.* **79,** 369 (1988).
18. G. L. BRETTHORST, J. J. KOTYK, AND J. J. H. ACKERMAN, *Magn. Reson. Med.* **9,** 282 (1989).
19. G. L. BRETTHORST, *J. Magn. Reson.* **88,** 533 (1990).
20. G. L. BRETTHORST, *J. Magn. Reson.* **88,** 552 (1990).
21. G. L. BRETTHORST, *J. Magn. Reson.* **88,** 571 (1990).
22. G. L. BRETTHORST, *J. Magn. Reson.* in press.
23. H. JEFFREYS, "Theory of Probability," Oxford Univ. Press, London, 1939: Later editions, 1948, 1961.
24. REV. T. BAYES, *Philos. Trans. R. Soc. London* **53,** 370 (1763): reprinted in *Biometrika* **45,** 293 (1958): "Facsimiles of Two Papers by Bayes," Hafner, New York, 1963.
25. WILLIAM H. PRESS, BRIAN P. FLANNERY, SAUL A. TEUKOLSKY, AND WILLIAM T. VETTERLING, "Numerical Recipes: The Art of Scientific Computing," Cambridge Univ. Press, Cambridge, 1989.
26. G. L. BRETTHORST, *in* "Lecture Notes in Statistics," (J. Berger, S. Fienberg, J. Gani, K. Krickeberg, and B. Singer, Eds.), Vol.**48,** Springer-Verlag, New York, 1988.