LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

LA-UR--89-3094

DE90 001835

CONF-8908106--2

OCT 1 1989

TITLE DYNAMICAL PROPERTIES FROM QUANTUM MONTE CARLO BY THE MAXIMUM ENTROPY METHOD

AUTHOR(S) R. N. Silver, T-11/LANSCE
D. S. Sivia, T-11/LANSCE
J. E. Gubernatis, T-11

SUBMITTED TO Workshop on Quantum Simulations of Condensed Matter Phenomena, Los Alamos, August, 1989
To be published by World Scientific Press

DISCLAIMER

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

FORM NO. 836 R4
ST. NO. 2629 5/81

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED MASTER

DYNAMICAL PROPERTIES FROM QUANTUM MONTE CARLO BY THE MAXIMUM ENTROPY METHOD

R. N. SILVER[1], D. S. SIVIA[1], J. E. GUBERNATIS[2]
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

ABSTRACT

An outstanding problem in the simulation of condensed matter phenomena is how to obtain dynamical information. We consider the numerical analytic continuation of imaginary time Quantum Monte Carlo data to obtain real frequency spectral functions. We suggest an image reconstruction approach which has been widely applied to data analysis in experimental research, the Maximum Entropy Method (MaxEnt). We report encouraging preliminary results for the Fano-Anderson model of an impurity state in a continuum. The incorporation of additional prior information, such as sum rules and asymptotic behavior, can be expected to significantly improve results. We also compare MaxEnt to alternative methods.

1. Introduction

The dynamical properties of strongly correlated many-body systems are of fundamental interest in most areas of condensed matter research. Quantum Monte Carlo methods [1] for the numerical simulation of such systems provide data on thermodynamic Green's functions in imaginary time. While this is appropriate for the determination of thermodynamic properties, it is difficult to extract the dynamical properties from the data because an analytic continuation from imaginary time to real time is required. The transform which relates the imaginary time data to the real frequency spectral function is similar to a Laplace transform. If one has physical reasons to believe in a model for the spectral function, the parameters of the model can be determined by fitting to the data. In the absence of a model, which is generally the case, the problem of infering the spectral function by inverting such transforms can be extremely ill-posed when the data are noisy and incomplete [2].

We suggest that this analytic continuation is essentially an image reconstruction problem, which is similar to others in a wide variety of experimental fields including radio astronomy, magnetic resonance imaging, photographic image en-

[1] MS B262, Theoretical Division and Los Alamos Neutron Scattering Center

[2] MS B262, Theoretical Division

hancement, neutron scatterir g, etc. The *image* sought in the present case is the real frequency spectral function. Almost any of the successful methods for image reconstruction should work with sufficiently good data. Given the limited data usually available, image reconstruction methods improve as they incorporate more prior knowledge about the quantity of interest. In choosing an image reconstruction method for our problem, the criteria to be considered should include: 1) the conceptual foundations of the statistical inference approach; 2) the efficiency of the method in making the most out of the prior knowledge and the data available, with the caution that the method should tend to put structure in the image only if there is statistically significant evidence for it in the data; and 3) the computational cost and programming effort required.

The Maximum Entropy Method (MaxEnt) [3] meets the first two criteria admirably, but it can score poorly on the third. For the analytic continuation problem, however, this weakness is of little consequence because the computational cost of MaxEnt is negligible compared to the cost of the Quantum Monte Carlo calculations required to produce the data. The programming effort required is also negligible because we use an existing third generation commercial code for MaxEnt image processing [4].

MaxEnt, like many successful image reconstruction techniques, can be understood in terms of Bayesian probability theory [5], also called the Law of Conditional Probabilities, which provides a logically consistent approach to statistical inference. Bayes' theorem encapsulates the learning process: our state of knowledge about some quantity of interest after an experiment (or simulation) depends on both the relevant data and our prior state of knowledge (or the lack thereof).

For the analytic continuation problem, the most important prior knowledge is that the spectral function is a positive additive probability distribution. Additional prior knowledge may include sum rules, symmetry properties, asymptotic behavior, etc. MaxEnt is a special case of Bayes' theorem appropriate to such probability distributions. It enforces the positivity of the spectral function, it tends to put structure in the image only if warranted by the data, and for most applications it has no adjustable parameters. When there are adjustable parameters, these can (in principle) also be estimated using Bayesian methods. Because MaxEnt is based on probability theory, it can provide error estimates of the reliability of the features in the spectral function given an adequate characterization of the statistical errors in the data. There is a vast body of experience in using MaxEnt on diverse image reconstruction problems [6], which provides insight into the images MaxEnt will produce for our problem.

Figure 1 shows an example of the use of the MaxEnt method for photographic image enhancement. The photograph of a "getaway" car is blurred due to

car motion observed with a finite camera shutter speed. The blurring function is evident from the streak produced by the point source to the right of the license plate. This is used as prior knowledge in a MaxEnt reconstruction. While the number on the license plate cannot be read from the original photo, it is easily read from the MaxEnt reconstruction.

Data **Reconstruction**

Figure 1
Optical deconvolution: an example of the use of MaxEnt.
(Courtesy of Drs. Gull & Skilling)

The present paper describes the application of MaxEnt to the analytic continuation problem for Monte Carlo data. We use the spectral function of the Fano–Anderson model of a bound state in a continuum for illustration [7]. We compare MaxEnt to other approaches to the analytic continuation problem in the recent literature by M. Jarrell and O. Biham [8], and by S. White et al. [9]. The problem has also been approached by H. Schuttler and D. Scalapino [10] using least square fitting, and by J. Hirsch [11] using Pade approximants.

2. The Analytic Continuation Problem

Quantum Monte Carlo generates the imaginary time (Matsubara) Greens

function

$$G(\tau) \equiv - \langle T_\tau[\hat{a}(\tau)\hat{a}^\dagger(0)] \rangle \quad , \qquad (1)$$

where τ is imaginary time, T_τ denotes time ordering, $\hat{a}^\dagger$ and $\hat{a}$ are particle creation and annhilation operators, and $\langle \ldots \rangle$ denotes a grand canonical ensemble average over the states of the many-body system. Knowledge of $G(\tau)$ in the range $0 \le \tau \le \beta$, where $\beta \equiv 1/k_B T$, is complete because $G(\tau)$ is periodic for bosons and antiperiodic for fermions. The dynamical properties of interest are given by the density of single-particle states which is termed the spectral function, $A(\omega)$. This is related to $G(\tau)$ by a transform

$$G(\tau) = \int_{-\infty}^{\infty} d\omega \, A(\omega) \, \frac{e^{-(\omega-\mu)\tau}}{1 + e^{-(\beta-\mu)\tau}} \quad , \qquad (2)$$

where μ is the chemical potential. The problem is to invert this transform in order to determine $A(\omega)$ from $G(\tau)$. Because this involves going from imaginary time to real time, it is equivalent to analytic continuation.

The Monte Carlo calculations provide data, $\{G_d(\tau_k)\}$ (i.e. calculated at a discrete set of τ points, $\{\tau_k\}$, distributed in the range $0 \le \tau_k \le \beta$), about $G(\tau)$ which are incomplete and noisy (i.e. subject to statistical error). In this case, the inversion is ill-posed because there exist an infinity of $A(\omega)$ all of which fit the data according to a chi-squared measure. Like the inversion of a Laplace transform, the analytic continuation problem is *extremely* ill-posed, because very small changes in the data can lead to very large changes in $A(\omega)$. The fitting of an $A(\omega)$ to the data by minimization of chi-squared (the method of least squares) is likely to put statistically insignificant structure into the $A(\omega)$ obtained. This non-uniqueness and extreme sensitivity to error have limited the ability to infer $A(\omega)$ from data produced by Quantum Monte Carlo calculations. The data may also be subject to systematic error due to the limitations of the Quantum Monte Carlo calculation, but we do not consider these further in this paper.

3. The Maximum Entropy Method

The image reconstruction approach we adopt is to infer the "best" $A(\omega)$ according to probability theory arguments. The method should tend to put structure

in $A(\omega)$ only if there is statistically significant evidence for it in the data. The method must also be capable of providing estimates of the statistical reliability of the results. The method should enforce whatever prior knowledge about $A(\omega)$ we may have, such as positivity and sum rules; i.e.

$$A(\omega) \geq 0 \quad , \qquad (3)$$

and

$$\int_{-\infty}^{\infty} d\omega\, A(\omega) = 1.0 \quad . \qquad (4)$$

Bayes' theorem provides our basis for statistical reasoning. It is simply the law of conditional probabilities:

$$P[X, Y] = P[X|Y] \times P[Y] = P[Y|X] \times P[X] \quad . \qquad (5)$$

Here, $P[X,Y]$ is the joint probability of X and Y, and $P[X|Y]$ is the conditional probability of X given Y. To specialize this to the analytic continuation problem, three quantities enter into the theorem. The first is the probability distribution of $A(\omega)$ before the experiment (or Monte Carlo calculation) is conducted, $P[A(\omega)]$, which is termed the *Prior* probability. For example, since we know that $A(\omega)$ cannot be negative, $P[A(\omega)]$ should be zero for $A(\omega)$ which go negative. The second is the conditional probability of producing the data, $\{G_d(\tau_k)\}$, via the experiment from a given $A(\omega)$, which is termed the *Likelihood* function, $P[G_d(\tau)|A(\omega)]$. It represents the modification of the Prior probability on $A(\omega)$ by the experiment. The third is the conditional probability of $A(\omega)$ given the data after the experiment, $P[A(\omega)|G_d(\tau)]$, which is termed the *Posterior* probability. The theorem states that the Posterior probability is proportional to the product of the Prior probability and the Likelihood function; i.e.

$$P[A(\omega)|G_d(\tau)] \propto P[G_d(\tau)|A(\omega)] \times P[A(\omega)] \quad . \qquad (6)$$

The *image* to be presented, $A_I(\omega)$, is the $A(\omega)$ which maximizes the Posterior probability. The statistical reliability of the image is to be obtained from the variation of $P[A(\omega)|G_d(\tau)]$ about this maximum.

The Likelihood function contains the new information provided by the experiment. If the data, $\{G_d(\tau_k)\}$, are assumed to be independent and Gaussian

distributed with error σ_k, the Likelihood function is related to the chi-squared measure by

$$P[G_d(\tau)|A(\omega)] \propto \exp\left(-\frac{\chi^2}{2}\right) \quad . \qquad (7)$$

where

$$\chi^2 = \sum_k \frac{[G_d(\tau_k) - G_f(\tau_k)]^2}{\sigma_k^2} \quad . \qquad (8)$$

Here, $G_f(\tau_k)$ are the data which a given choice of $A(\omega)$ would produce in the absence of noise. For example, if the $A(\omega)$ is to be determined in pixels at discrete frequencies, $\{\omega_i\}$, then

$$G_f(\tau_k) = \sum_i A(\omega_i)\Delta\omega \frac{e^{-(\omega_i-\mu)\tau_k}}{1+e^{-(\omega_i-\mu)\beta}} \quad . \qquad (9)$$

If one ignores the Prior probability in Eq. (6), i.e. $P[A(\omega)] \rightarrow constant$, then maximizing the Posterior probability is equivalent to fitting the data by the method of least squares, i.e. minimizing χ^2.

The errors will be assumed to be independent in all the simulations of this paper. It is straightforward to generalize the Likelihood function to non-independent data. This may be important for real Monte Carlo simulations since there is nothing in the logic of the Monte Carlo method to assure statistical independence.

The Maximum Entropy Method corresponds to a particular choice of Prior which incorporates the prior knowledge that the spectral function is positive and additive. In that case, a variety of different statistical inference arguments [3,12] lead to the conclusion that the Prior should take the form

$$P[A(\omega)] \propto \exp(\alpha S) \quad . \qquad (10)$$

Here, S is the information theory entropy [13] of the image, i.e.

$$S = \sum_i \Delta\omega \left[A(\omega_i) - m(\omega_i) - A(\omega_i)\ln\left(\frac{A(\omega_i)}{m(\omega_i)}\right)\right] \quad . \qquad (11)$$

S is defined relative to a starting model, $m(\omega)$, for $A(\omega)$. It is termed the *default model*, because it is the image to which the Maximum Entropy Method will default in the absence of data. It is usually chosen to be the smoothest function consistent

with prior knowledge such as sum rules. If, before the Quantum Monte Carlo simulation, we are completely ignorant about $A(\omega)$ except for the sum rule, we take $m(\omega)$ to be flat with a magnitude to satisfy the sum rule in the frequency range of interest. More generally, $m(\omega)$ should be chosen by the maximum entropy method using the prior knowledge (e.g. moments) as constraints (or "testable information" [3]).

While space does not permit a detailed presentation of the variety of rationales leading to Eqs. (10-11), condensed matter scientists will understand a statistical mechanics analogy. The thermodynamic entropy is the logarithm of the number of different states by which one can arrive at a given total energy, or any other macroscopic constraint. Analogously, the information theory entropy is the logarithm of the number of ways by which one can arrive at a given $A(\omega)$ in a Poisson process from the default model. For a flat $m(\omega)$, this is called the "monkey argument" by the image processing community [14]:

> *If a team of monkeys throws a very large number* N *of quanta randomly at the* M *a priori equivalent cells of an image, then the probability of obtaining a particular set* $(n_1, n_2, \cdots, n_M)$ *of occupation numbers shall be proportional to the degeneracy* $N!/n_1!n_2! \cdots n_M!$.

More generally, in a Poisson process the expected number of counts in a cell (or pixel) is $\bar{n}_i = \alpha\, m(\omega_i)\Delta\omega$. The parameter α characterizes the degree of fluctuation about the default model. Application of Stirling's formula to the factorials in a Poisson probability distribution then leads to Eqs. (10-11).

Finally, the image is obtained by maximizing the Posterior probability. This is the same as functional variation of $A(\omega)$ to maximize the entropy, S, subject to a constraint on χ^2 with Lagrange multiplier $1/\alpha$; hence the name *Maximum Entropy Method*. Clearly, $1/\alpha$ controls the relative importance of the prior knowledge and the data. It is a statistical regularization parameter which is fixed by self-consistency arguments. Before the experiment the image starts at the default model and $\alpha = \infty$. As the data improve α becomes smaller, and the image deviates from the default model (hopefully) toward the correct $A(\omega)$ acquiring sharper resolution. In traditional (or *Historic*) MaxEnt, $1/\alpha$ is chosen so that the image maximizes the entropy subject to it being "feasible", in that it fits the data according to $\chi^2 \leq N_r$, where N_r is the number of independent data. This is, of course, consistent with

the χ^2 - probability distribution for independent Gaussian errors. In the more modern (or *Classic*) MaxEnt [15], $1/\alpha$ is itself determined by Bayesian arguments which maximize the Posterior probability of α given the data, $P[\alpha|G_d(\tau)]$. The image exhibited is the center of a probability distribution (called the "bubble") of possible images. It is the Classic MaxEnt formulation that makes possible estimates of the statistical significance of the features in the image.

Maximum Entropy is a statistical inference principle, how one implements the principle is an algorithm, choosing the appropriate prior knowledge is physics, and incorporating it in the algorithm is an art. Prior knowledge should be built into the default model and into the data as a constraint. The images are conditional on the prior knowledge. The use of prior knowledge can be essential to solving some problems, but invalid prior knowledge can also lead to spurious results. For example, using a functional model for $A(\omega)$ defined by a few parameters is a strong form of incorporating prior knowledge. However, if the model cannot be fit to the data, at least some of the "prior knowledge" contained in the model is wrong. In many image processing situations, one uses the flexibility of MaxEnt to iteratively input increasing amounts of prior knowledge (e.g. first positivity only, then include symmetry, and so on.), or to test the validity of assumed "prior knowledge". Often, an initial MaxEnt image will suggest additional prior knowledge which should be incorporated. MaxEnt images often lead to a physical model for the quantity of interest, the parameters of which can then be estimated directly using the raw data.

4. Examples of MaxEnt Analytic Continuation

Let us consider the spectral function shown in Fig. 2(a) termed the "truth", where we have binned the spectral function into 41 pixels in the range $-\Omega \le \omega - \mu \le \Omega$. Using this in the transform, Eq. (2), with $\beta = 10.0$, $\Omega = 5.0$, binning with $\Delta\tau = 0.125$, and adding 1% relative Gaussian random noise yields the $G(\tau)$ data shown in Fig. 2(b) termed the "mock data". The problem then is to recover the spectral function from this mock data. Fig. 2(c) shows the MaxEnt reconstruction of the spectral function using the Classic MaxEnt algorithm. One can see that the basic structure of the truth is recovered in the reconstruction. However, the features at large $|\omega - \mu|$ are broadened compared to features at small $|\omega - \mu|$.

Most significantly, Classic MaxEnt allows one to place error bars on the

integrated intensities of the features in the image. The areas to be integrated over are marked by arrows. The corresponding integrated intensity and estimated error are shown. The central peak and the gap for small $|\omega-\mu|$ are well determined, but the side peaks and the gap at large $|\omega-\mu|$ are less well determined.

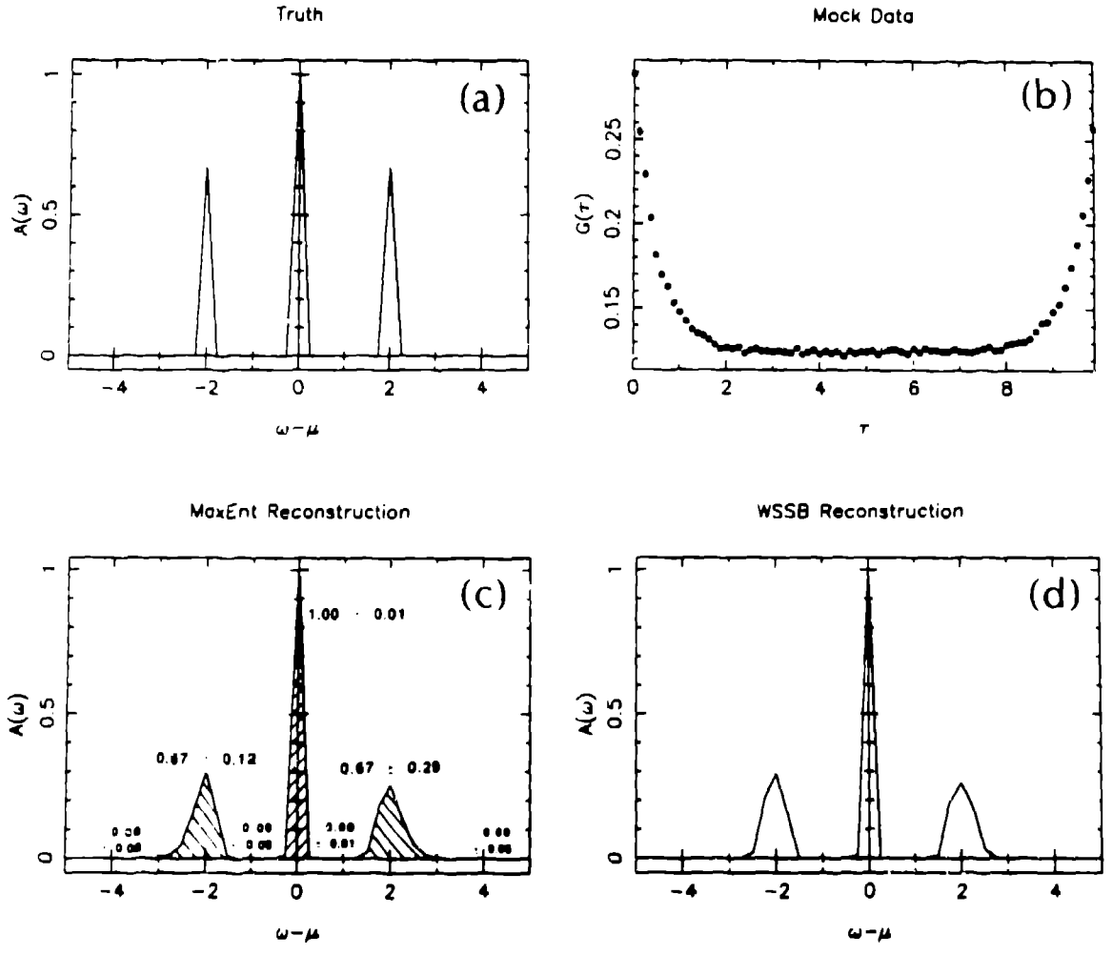

Figure 2 **First simulation for a spectral function with only sharp features, with parameters Ω=5, $\Delta\omega$=0.25, β=10, and $\Delta\tau$ =.125: a) true spectral function; b) resultant mock data, $G(\tau)$, with relative Gaussian random error of 1%; c) MaxEnt reconstruction; d) reconstruction using the WSSB [9] algorithm.**

We compare MaxEnt with the method of S. White, et al. (WSSB) [9]. Their method is to minimize a modified least squares measure for the data which in our notation is

$$\chi^2_{WSSB} = \sum_k \frac{[G_d(\tau_k) - G_f(\tau_k)]^2}{\sigma_k^2} + b \sum_i (A(\omega_i) - A(\omega_{i+1}))^2 + h \sum_i \theta(-A(\omega_i))A^2(\omega_i) \quad . (12)$$

The first term is conventional. The second term with parameter b is designed to encourage smoothness and, hence, minimize the structure in the reconstruction. The third term with parameter h is designed to encourage (but not enforce) the positivity of the spectral function. The authors recommend how to iteratively adjust the parameters b and h to arrive at $\chi^2 = 1.1\ \chi^2_{min}$, where χ^2_{min} is the lowest value attained when enforcing positivity but very little smoothness. The resulting image is shown in Figure 2(d). This image is essentially the same as the MaxEnt image, but it lacks information about the statistical reliability of the image. The WSSB approach is sensible since it incorporates the additional prior knowledge of positivity, but it is also *ad hoc* because there is no self consistent way of choosing the parameters b and h. Moreover, the smoothness assumption (correlations between neighboring pixels) is invalid prior knowledge, as there are spectral functions which are not smooth (such as the Fano–Anderson model we consider later).

The pathology of the smoothness assumption can be shown with the spectral function in Fig. 3(a) which includes a sharp peak at $|\omega - \mu| = 0$ and broad structure at large $|\omega - \mu|$. The parameters of the simulation are the same as in Fig. 2. Figure 3(b) shows the corresponding mock data. Figure 3(c) shows the WSSB reconstruction at $\chi^2 = 1.1\ \chi^2_{min}$ which produces spurious structure. Figure 3(d) shows the image obtained by turning the smoothness up to achieve $\chi^2 = 2.0\ \chi^2_{min}$, and Figure 3(e) has further smoothing with $\chi^2 = 10.0\ \chi^2_{min}$. One can see that the fits to the broad structure are improving at the expense of broadening the central peak. Finally, Fig. 3(f) shows the MaxEnt reconstruction (symmetry was not enforced). It is much closer to the original image, and it permits estimates of the statistical reliability of features in the image, as shown.

The method of Jarrell and Biham [8] also attempts to enforce smoothness and positivity while working with data in Matsubara frequency space rather than imaginary time. There is no explicit provision for statistical error propagation in their method. It shows similar qualitative behavior to the WSSB method.

For a more realistic example of MaxEnt analytic continuation, we consider the spectral function of the Fano–Anderson model of an impurity state in a continuum. Here, the Hamiltonian is

$$\hat{H} = e_c \hat{b}^\dagger \hat{b} + \sum_k e_k \hat{a}_k^\dagger \hat{a}_k + \sum_k A_k (\hat{a}_k^\dagger \hat{b} + \hat{b}^\dagger \hat{a}_k) \quad , \qquad (13)$$

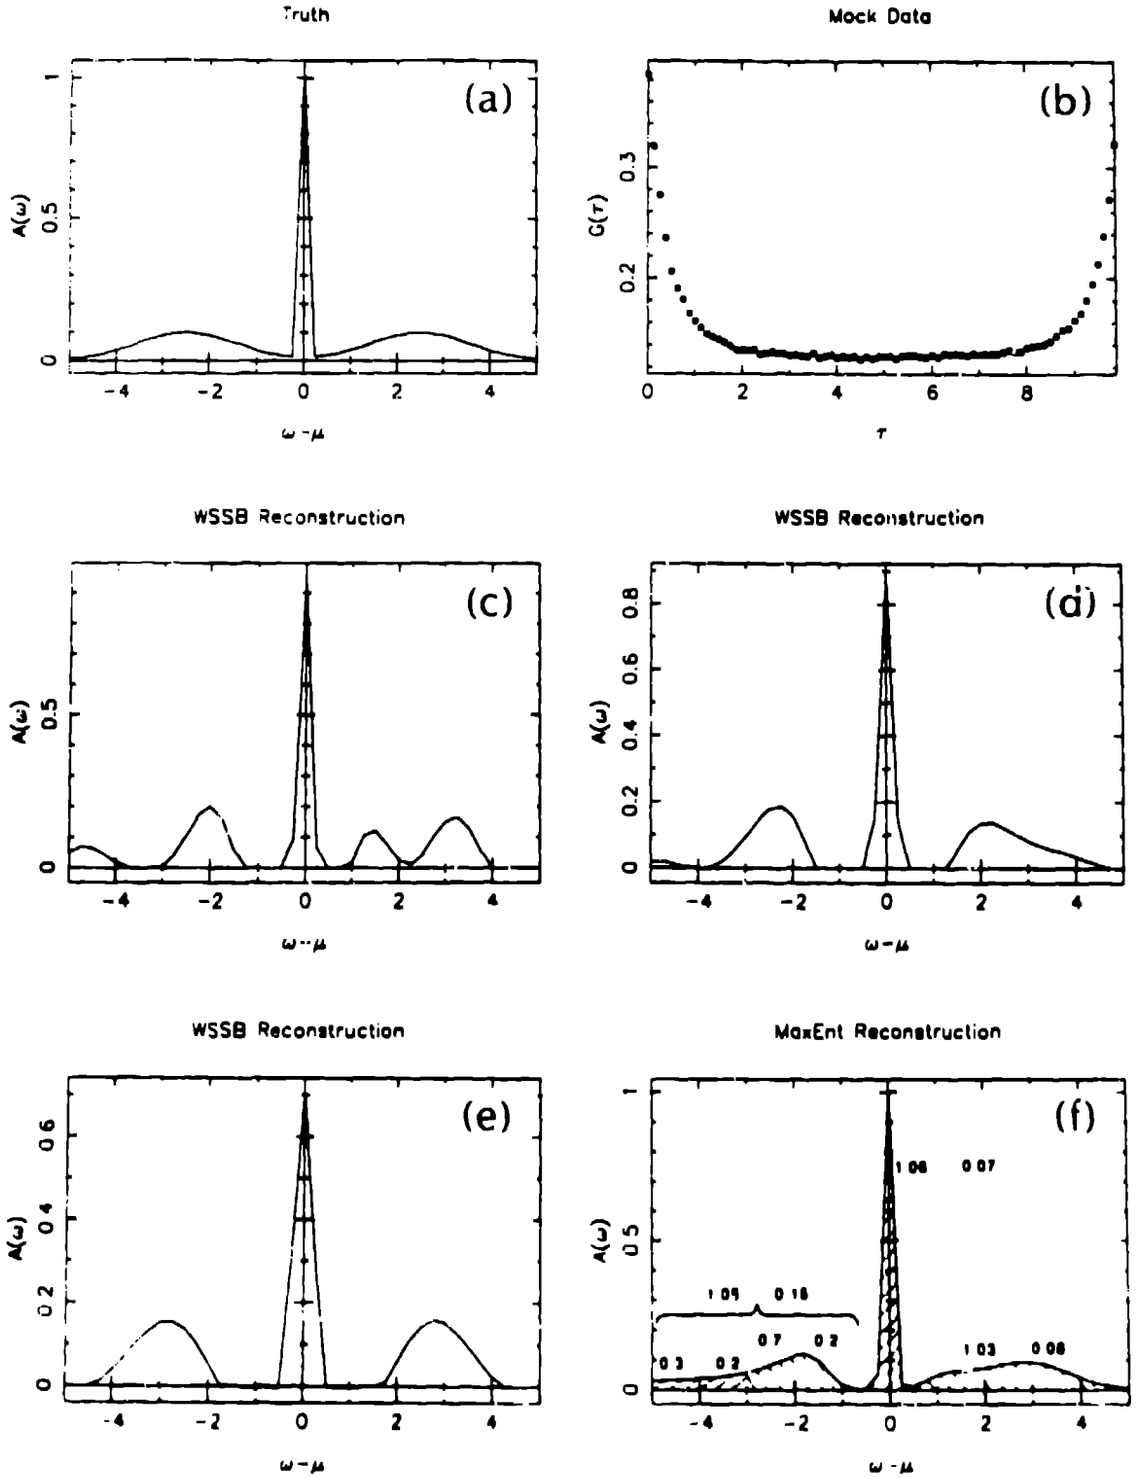

Figure 3 Second simulation for a spectral function having both sharp and smooth features, as explained in the text.

where we consider a one-dimensional band with

$$A_k = \sqrt{\frac{c}{L}} \quad ; \quad e_k = -w\cos k \quad . \qquad (14)$$

Here $\hat{b}^\dagger$ is the creation operator for an electron on the impurity site, $\hat{a}_k^\dagger$ is the creation operator for an electron in the band, e_c is the impurity level energy, $2w$ is the bandwidth, A_k is the coupling between the impurity level and the band, and L is the lattice spacing. For this model, one can analytically derive the spectral function for the impurity state from the imaginary time impurity Greens function

$$G_b(\tau) = - < T_\tau [\hat{b}(\tau)\hat{b}^\dagger(0)] > \quad . \qquad (15)$$

For a variety of choice of parameters, we can generate mock $G_b(\tau)$ data by discretizing in τ and adding Gaussian random noise. We then examine the ability of MaxEnt to recover the impurity spectral function from such data.

Figure 4 shows the resulting simulations with parameters as indicated in the figures. In each case, the left-most figure is the spectral function for the indicated choice of parameters, the middle figure are the mock data for $G(\tau)$, and the right-most figure is the MaxEnt reconstruction. The total area under the spectral function is one. The continuum is always chosen to extend from -1.0 to 1.0, but the position of the impurity level and its coupling to the continuum is varied. Discrete states are split off from the continuum which contribute delta functions to the spectral functions (indicated by the solid vertical lines) of weights Z_+ and Z_-. N_τ is the number of data points of $G(\tau)$; *Error* is the fractional standard deviation of the error in $G(\tau)$; β is the inverse temperature taken to be 30.0; and μ is the chemical potential taken to be 0.0. The *Chi Sq* is the actual chi-squared achieved in the MaxEnt reconstruction, to be compared with the N_τ. Qualitative and semi-quantitative featues of the true spectral function are reproduced by the MaxEnt reconstruction, but there is a tendency toward broadening at large $|\omega|$ and toward cuspy structure at small $|\omega|$. The error analysis of the Classic MaxEnt method indicates that the cuspy structure is not statistically significant. The $G_b(\tau)$ data generated now has many more τ points and much greater statistical accuracy than in the simulations of Figs. 2 and 3 in order to provide sufficient information to reconstruct the more complex spectral function. These errors are smaller than those used in refs. [8,9,11], but they are within the range of the current state-of-the-art in Quantum Monte Carlo simulations. The data shown in the middle of the figure are clearly insensitive to the detailed structure in the spectral function.

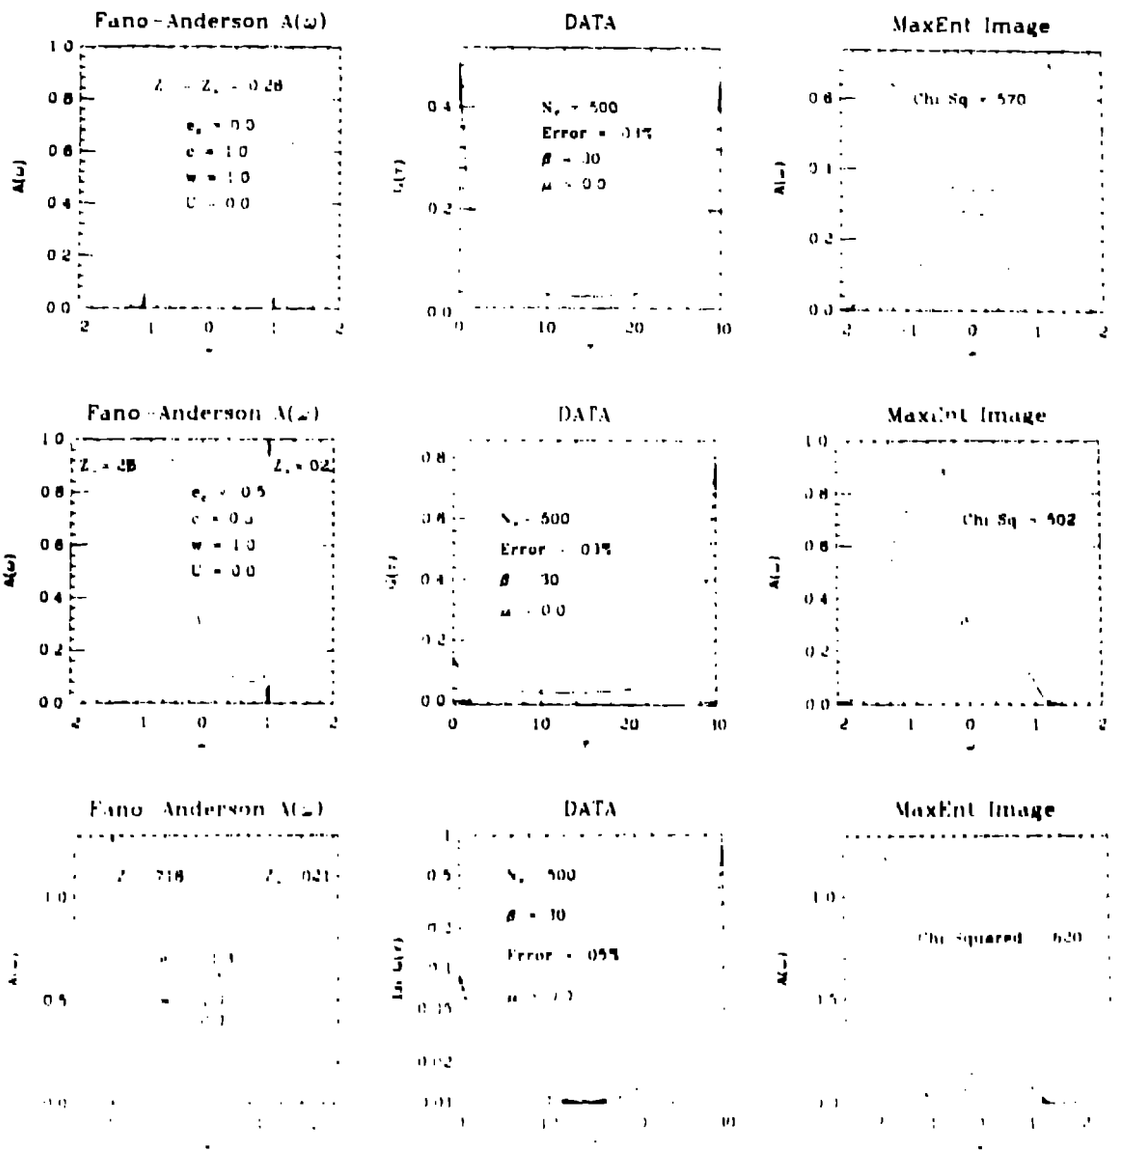

Figure 4 MaxEnt simulations for the Fano-Anderson model of an impurity state in a continuum.

which is what makes this problem extremely ill-posed. Nevertheless, one can see that the MaxEnt images on the right of the figure have recovered the essential features of the original spectral functions shown on the left. However, there is a tendency to put cuspy structure in the image for small $|\omega - \mu|$, and to broaden structure at large $|\omega - \mu|$. One can show that this tendency is a feature of any image reconstruction method. The cuspy structure at small $|\omega - \mu|$ can be shown to be spurious by again using the Classic MaxEnt option of estimating the statistical errors on the integrated intensity of that structure.

5. Conclusions

Using an image reconstruction approach, the Maximum Entropy Method, we have shown that imaginary time Quantum Monte Carlo data can be analytically continued to obtain real frequency spectral functions. We have reported encouraging preliminary results for the spectral function of the Fano-Anderson model of an impurity state coupled to a continuum. We have compared our results to other proposals in the literature. This success opens the door to obtaining dynamical properties of quantum many body systems by Quantum Monte Carlo methods [16].

A more detailed paper [17] presents a Likelihood function analysis of statistical error propagation (and ill-posedness) in the analytic continuation problem, which is independent of the choice of image reconstruction method. This includes an explanation of the tendency to put sharp structure at small $|\omega - \mu|$ and broad structure at large $|\omega - \mu|$, the optimization of Quantum Monte Carlo simulations, and the incorporation of covariance of statistical error propagation in Quantum Monte Carlo calculations.

6. Acknowledgements

R. N. S. and D. S. S. gratefully acknowledge the support of the Office of Basic Energy Sciences / Division of Materials Science of the U. S. D. O. E. for the Los Alamos Neutron Scattering Center. J. E. G. gratefully acknowledges support from the Energy Research Supercomputing Access Program of the U. S. D. O. E.

7. References

1. For a tutorial review of Fermion Monte Carlo methods, see E. Loh, J. E. Gubernatis, "Stable Numerical Simulations of Models of Interacting Electrons in Condensed Matter Physics", in *Electronic Phase Transitions*, ed. W. Hanke and

Y. Kopaev, Elsevier, New York (1990).
2. F. S. Acton, *Numerical Methods that Work*, Harper and Row Publishers, New York (1970); see the section on p. 245–257: "Interlude: What *Not* to Compute".
3. E. T. Jaynes, "Where Do We Stand on Maximum Entropy?" (1978), reprinted in *E. T. Jaynes, Papers on Probability, Statistics, and Statistical Physics*, ed. R. Rosenkrantz, (1983), Reidel Publishing Company.
4. MEMSYS 3, produced by Maximum Entropy Data Consultants Ltd. (1989); this is an implementation of the Cambridge algorithm. For details of MEMSYS 1, see J. Skilling, R. K. Bryan, Monthly Notices of the Royal Astronomical Society, **211**, 114 (1984).
5. H. Jeffreys, *Theory of Probability*, Oxford University Press, 1939; R. P. Cox, Am. J. Phys. **17**, 1–13 (1946).
6. S. F. Gull, J. Skilling, IEE Proc. **131**(F), 646 (1984); R. Narayan, R. Nityananda, Ann. Rev. Astrophys. **24**, 127 (1986).
7. G. D. Mahan, *Many Particle Physics*, Plenum Press, New York (1981), p. 256ff.
8. M. Jarrell, O. Bihnm, to be published.
9. S. R. White, D. J. Scalapino, R. L. Sugar, N. E. Bickers, to be published.
10. H. B. Schuttler, D. J. Scalapino, Phys. Rev. Lett. **55**, 1204 (1985); Phys. Rev. B**34**, 4744 (1986).
11. J. Hirsch, "Simulation of Magnetic Impurities in Metals", in *Quantum Monte Carlo Methods*, ed. M. Suzuki, Springer-Verlag, Heidelberg (1987).
12. J. E. Shore, R. W. Johnson, IEEE Trans. Info. Theory **IT-26**, 26 (1980); ibid, **IT-29**, 942 (1980); Y. Tikochinsky, N.Z. Tishby, R. D. Levine, Phys. Rev. Lett. **52**, 1357 (1984).
13. C. E. Shannon, Bell Systems Tech. J. **27**, 379–423, 623–656 (1948).
14. S. F. Gull, G. J. Daniell, Nature **272**, 686 (1978).
15. J. Skilling, "Classic Maximum Entropy", in *Maximum Entropy and Bayesian Methods*, ed. J. Skilling, published by Kluwer Academic Publishers (1989).
16. See also the contributions of W. Miller et al. and Haug et al. in this volume.
17. R. N. Silver, D. S. Sivia, J. E. Gubernatis, submitted to Phys. Rev. B.