Techniques for Observing Solar Oscillations

Timothy M. Brown

High Altitude Observatory/National Center for Atmospheric Research*
P.O. Box 3000
Boulder, CO 80307
USA

*The National Center for Atmospheric Research is sponsored by the National Science Foundation.

1 ANALYSIS TOOLS AND THE SOLAR NOISE BACKGROUND

When we observe solar oscillations, we are concerned with measuring perturbations on the Sun that are almost periodic in space and time. The periodic waves that interest us are, however, embedded in a background of broadband noise from convection and other solar processes, which tend to obscure and confuse the information we want. Also (and worse), the "almost-periodic" nature of the waves leads to problems in the interpretation of the time series that we measure. Much of the subject of observational helioseismology is thus concerned with ways to minimize these difficulties.

1.1 Fourier Transforms and Statistics

A common thread runs through all of the analysis tricks that one plays when looking at solar oscillations data, and indeed through many of the purely instrumental concerns as well: this thread is the Fourier transform. The reason for this commonality is, of course, that we are dealing with (almost) periodic phenomena – either the acoustic-gravity waves themselves, or the light waves that bring us news of them. Since many of the same notions will recur repeatedly, it is worth taking a little time (and boring the cognoscente) to review some of the most useful properties of Fourier transforms and power spectra. In what follows, I shall simply state results and indicate some of the more useful consequences. We shall see below that even when the Big Theorems of Fourier transforms do not apply, (as with Legendre transforms, for instance), analogous things happen, so that the Fourier example is a helpful guide to the kind of problems we may have. Those who want proofs and details can find them in Bracewell's (1965) classic book.

1.1.1 Shift and Convolution Theorems

Suppose we have a (possibly complex) function $f(x)$. Its Fourier transform $F(k)$ is defined as

$$F(k) = \int_{-\infty}^{\infty} f(x) exp(-2\pi i k x)\, dx \quad . \tag{1}$$

To invert the transform, just invert the sign in the exponential:

$$f(x) = \int_{-\infty}^{\infty} F(k) exp(+2\pi i k x)\, dx \quad . \tag{2}$$

A useful notation is to denote transforms (either forward or inverse) by overbars on the function being transformed:

$$\overline{f(x)} = F(k) = \int_{-\infty}^{\infty} f(x) exp(-2\pi i k x)\, dx \quad . \tag{3}$$

Aside from a sign change in the exponential, the Fourier transform is its own inverse, so that (with the looseness provided by the overbar notation) we can write

$$\overline{\overline{f(x)}} = \overline{F(k)} = f(x) \, . \tag{4}$$

This feature is extremely helpful, since it means that many of the following relations between a function and its transform may be applied in either direction. Another important property of Fourier transforms is their scaling property:

$$\overline{f(ax)} = \frac{1}{a} F(\frac{k}{a}) \, . \tag{5}$$

A few special cases are worth noting: If $f(x)$ is real, then $F(k)$ is conjugate-symmetric in k, that is, $F(-k) = F^*(k)$. If $f(x)$ is real and symmetric in x, then $F(k)$ is purely real. If $f(x)$ is real and antisymmetric in x, then $F(k)$ is purely imaginary.

Straightforward computation shows that $F(k)$ obeys the *shift theorem*, namely that if $f(x)$ has the Fourier transform $F(k)$, then $f(x-a)$ has the transform $exp(-2\pi i a k) F(k)$.

The shift theorem may in turn be used to prove the single most useful result about Fourier transforms, the *convolution theorem*. A convolution $h(x)$ is defined as

$$h(x) = f \otimes g = \int_{-\infty}^{\infty} f(u) g(x-u) du \quad . \tag{6}$$

If the Fourier transforms of f and g are $\overline{f}$ and $\overline{g}$, then the convolution theorem says that

$$\overline{h} = \overline{f \otimes g} = \overline{f}\,\overline{g}. \tag{7}$$

That is, the Fourier transform of a convolution of two functions is the product of their Fourier transforms. Because if the self-inverse nature of the transform, it also follows that the transform of a product of two functions is the convolution of their transforms:

$$\overline{fg} = \overline{f} \otimes \overline{g} \, . \tag{8}$$

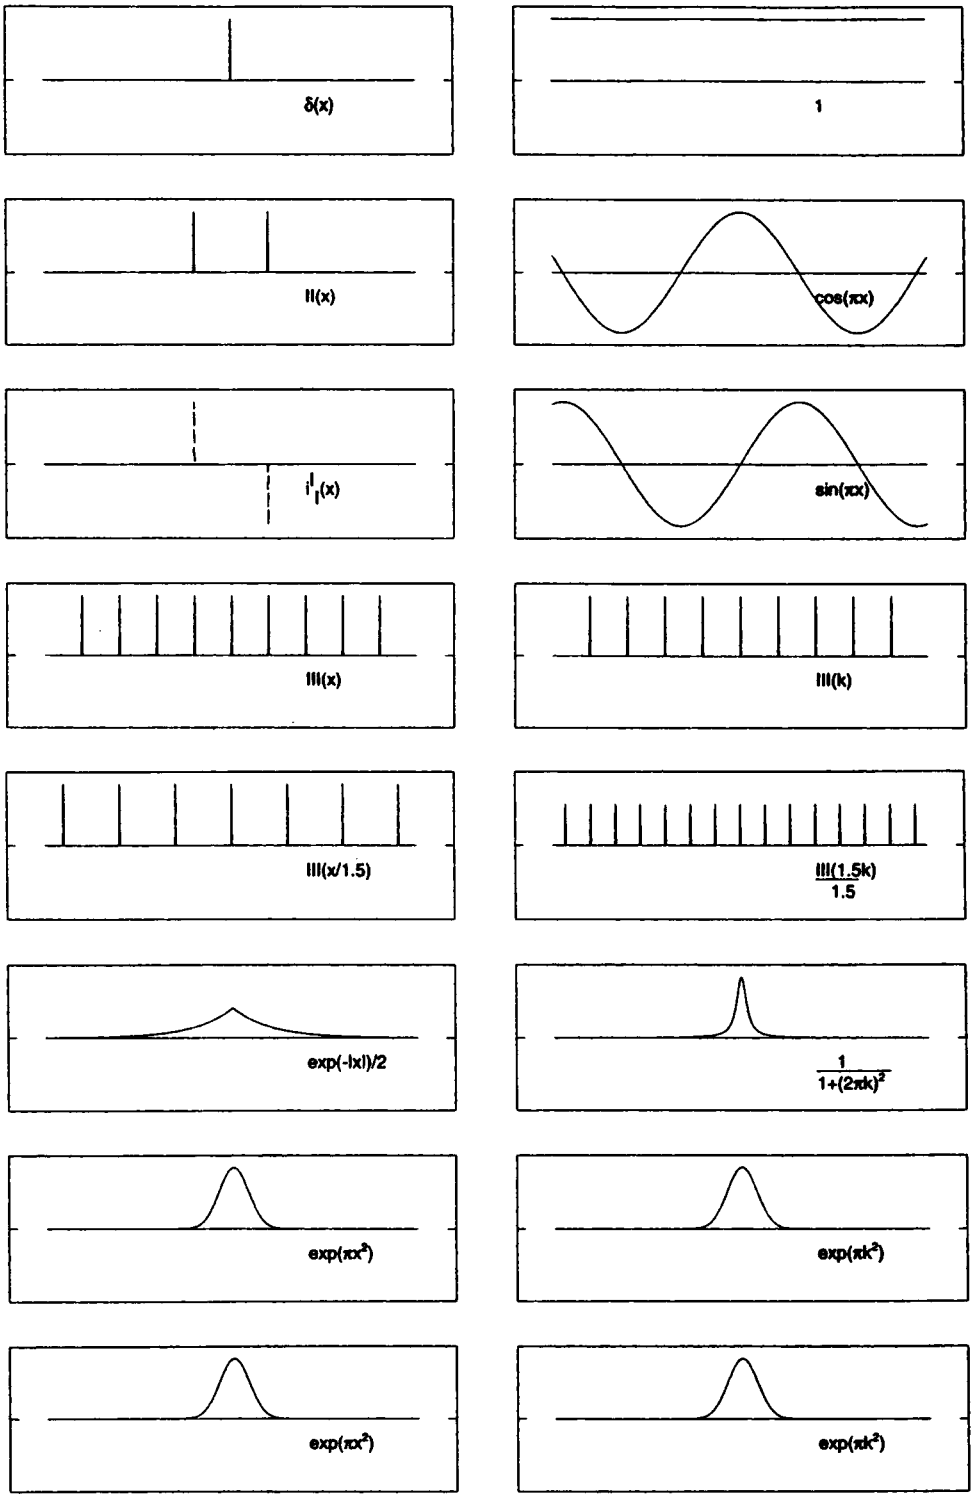

Figure 1: *Useful Fourier transform pairs (see Bracewell, 1965).*

1.1.2 Transforms of Useful Functions

At this point, it is useful to display a few commonly-encountered Fourier transform pairs. Those functions illustrated in fig. 1 will serve our needs for the rest of this discussion, but more may be found at the end of Bracewell's (1965) book. The one function illustrated here that may be unfamiliar is what Bracewell terms the "shah" function $III(x)$. It consists of an infinite string of equal delta functions with unit separation between them. This function is handy in many ways, not least being that it is its own Fourier transform. Notice that because of the scaling relation for Fourier transforms, spreading the delta functions farther apart in x causes the delta functions in k-space to grow closer together.

1.1.3 A Physical Example

To illustrate the use of some of these notions, let us consider a Fabry-Perot interferometer. The Fabry-Perot is doubly instructive in our context: it has been used to observe solar p-modes, and moreover it is a rough analogue for the cavity in which the p-modes propagate.

A Fabry-Perot interferometer consists of two plane mirrors lying parallel to one another, separated by a distance s. I will suppose that the mirrors reflect most of the light incident on them (independent of wavelength), and that the rest is transmitted. One way to understand how such an interferometer behaves is to imagine that a light source to the left of the interferometer emits a very short pulse of light in the form of a plane wave traveling toward the mirrors, with its propagation vector inclined at an angle θ to the normal to the mirrors. When this pulse hits the first mirror, some of the light is transmitted into the cavity between the mirrors. Part of this light is transmitted through the second mirror, and appears to an outside observer as an impulse at a time t_0. Of the light reflected from the second mirror, most is reflected again from the first mirror and arrives back at the second mirror after a time $\delta t = s/\cos(\theta)c$. Part of this delayed pulse is again transmitted by the second mirror, giving rise to a second output pulse at $t_0 + \delta t$. Subsequent reflections lead to a long string of pulses spaced δt apart in time, *i.e.*, a shah function in time. The frequency content of this output signal is given by the Fourier transform, namely a shah function in frequency, with peaks that are equally separated in frequency, with a spacing equal to $1/\delta t$ (see fig. 2a and 2b). These frequencies are the *resonant frequencies* of the interferometer; they evidently depend on the spacing s between the plates, the speed of wave propagation c, and on the propagation direction θ.

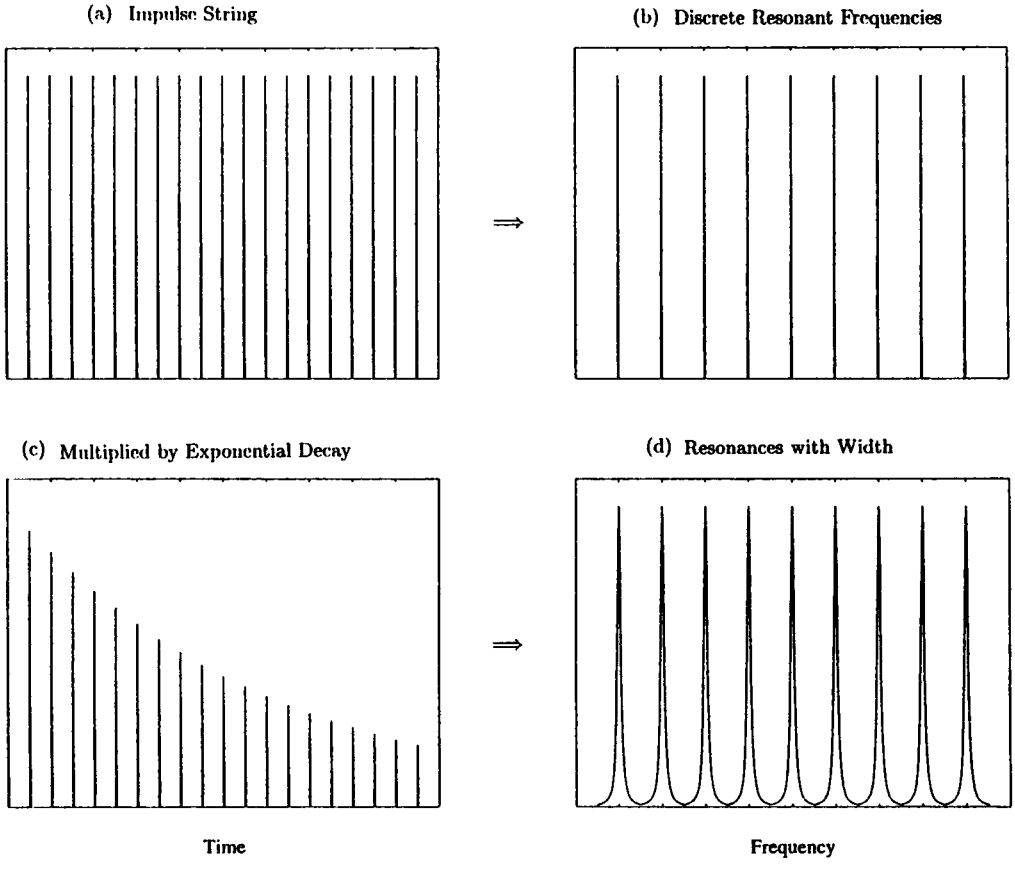

Figure 2: *(a) and (c): Time series of impulses emerging from a Fabry-Perot interferometer. (b) and (d): Corresponding Fourier transforms.*

An alternative way to analyze the Fabry-Perot interferometer is to insist that for resonant frequencies, an integral number of wavelengths of light must fit into one forward-and-back pass through the interferometer (with due allowance for phase shifts upon reflection, if any). This approach is equivalent to the travel-time argument just given; both are useful, with the clearest approach being dictated by the details of the situation. Here I emphasize the travel-time description mostly because it illustrates a general property of linear systems: The frequency response may be determined from the response of the system to a delta-function excitation.

Many physical systems (including the transmission of interferometers and the small-amplitude pulsations of stars) behave like this:

$$\text{Observed Quantity} = \Lambda(\text{Forcing Function}) , \tag{9}$$

where Λ is some linear operator that gives zero if the forcing function is always zero. (For instance, Λ may contain scaling factors, derivatives, or time delays.) Then by virtue of linearity and because of the orthogonality of the sines and cosines, the frequency components of any arbitrary forcing function may be considered one at a time; the effect of Λ is simply to multiply each frequency component by a (frequency dependent) complex factor. Since a δ function contains all frequency components in equal measure, the output resulting from a δ function input (the so-called *impulse response*) contains information about all frequencies. This is the "police brutality" view of systems analysis: if you want to find out all about a linear system, hit it once, and listen to what it has to say.

Now the Fabry-Perot analysis so far is unrealistic, because I have supposed that the string of pulses resulting from an impulse excitation will go on forever. In real life, the transmitted pulses bleed energy out of the cavity, so that the output impulse amplitude decreases exponentially with time (fig. 2c). What is the effect of this upon the spectrum? The new output signal is now the original shah function multiplied by an exponential envelope. By the convolution theorem, its transform is therefore the convolution of the transforms of the shah function (which is another shah function) and the transform of a decaying exponential (which is a Lorentzian). The effect is to change the resonances of the interferometer from infinitely sharp resonances to Lorentzian-shaped lines, whose widths are inversely proportional to the exponential decay time (fig. 2d). This sort of behavior turns out to be a pretty fair model of real interferometers, and is at least evocative of the processes that characterize the Sun's p-modes.

1.1.4 Stochastic Excitation

Suppose one has a system (like the Fabry-Perot) that is characterized by the spectrum of its impulse response, but one excites this resonator with a much more complicated driving function (a flock of impulses, say, with random amplitudes and times of occurrence). What is the resulting Fourier transform? Because of the system's linearity, the output signal (in the time domain) is the convolution of the system's impulse response with the driving function. The transform is then the product of the transforms of the impulse response and of the driving function. Many kinds of random excitation function have transforms whose expectation values are slowly-changing functions of frequency. However, a single realization of such a process usually looks very noisy (for the probability distributions related to such spectra, see the next section). In these cases, the emergent spectrum looks like noise that has been multiplied by the spectrum of the impulse response. In this sense, linear resonant systems act like filters, which

(in the frequency domain) multiply but do not mix the frequency components of their driving function.

1.1.5 Probability Distributions for Power Spectra

For very many naturally-occurring time series, it turns out that different Fourier transform components behave as independent random variables, with the real and imaginary parts each being normally distributed with identical variance (Groth, 1975). The power is the sum of the squares of the real and imaginary parts of each frequency component, so its probability density is not Gaussian, but rather exponential (also called a χ^2 distribution with 2 degrees of freedom):

$$P(s) = e^{-s} . \tag{10}$$

An important property of this probability distribution is that the standard deviation associated with it is equal to its expectation value. Thus, the difference between the typical power at a given frequency and the value actually measured in a single realization is usually about as large as the typical power itself. Near resonances where the mean power is large, the measurement errors will be large. Far from resonances, where little power is expected, the distribution of measured power values is relatively narrow. Moreover, the high-power wing of the distribution in eq. (10) extends much further than does a Gaussian with the same variance. This means that one sees large-power events much more often than would naively be expected.

These peculiarities of the probability distribution for power spectra carry with them a severe danger of over-interpretation. It is all too easy to see a big feature in a power spectrum and think that it must be real, when in fact it is merely a (not very unlikely) statistical fluke. This sort of thing is illustrated in fig. 3, which shows a number of power spectra of a stochastically-excited damped oscillator. All of the spectra are for the same oscillator; all that is different is the random forcing. In all cases the limit spectrum (the spectrum that would be obtained by averaging many realizations) is shown as a dashed line, and the actual spectrum for each realization is solid. Fig. 3a shows the average of 60 spectra, which does indeed approach the limit spectrum. Among the individual realizations, almost any desired pathology may be found: narrow lines, wide lines, lines displaced either left or right from the average line center, as well as clear and symmetric doublets and triplets. Obviously great care must be taken when interpreting such spectra if one is to avoid erroneous conclusions.

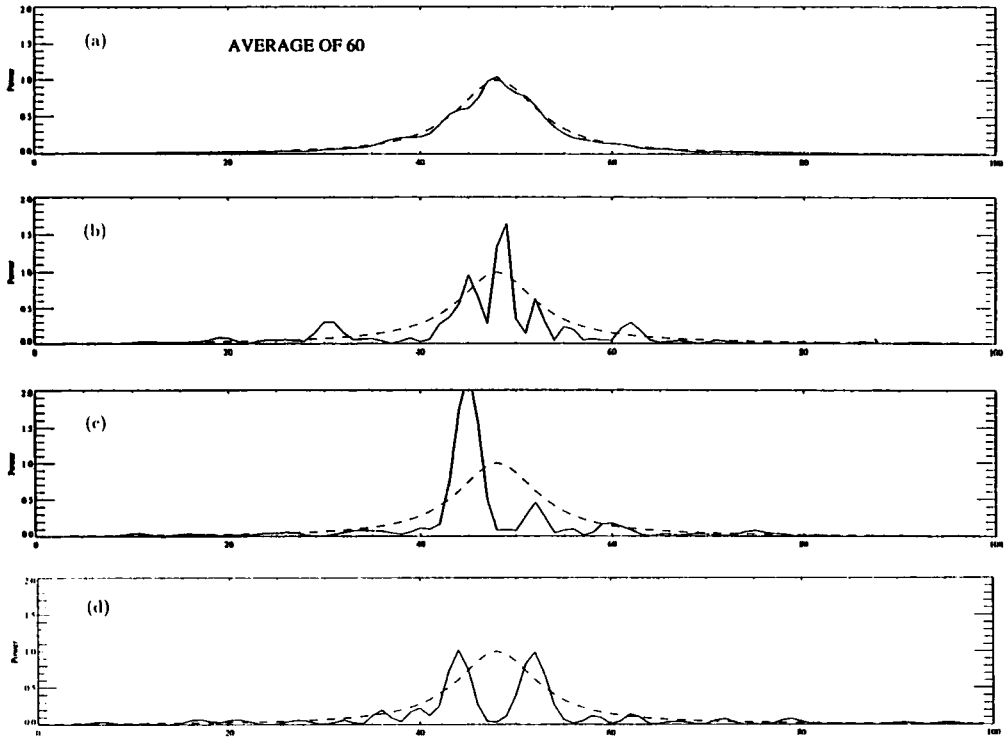

Figure 3: *Power spectra of a stochastically-driven damped oscillator. (a) shows the average of 60 realizations with different random driving. (b-d) show three single realizations, illustrating the wide range of behavior that characterizes such oscillators. In all cases, the limit spectrum is given by the dashed line.*

1.2 The Solar Noise Background

In addition to confusion resulting from random forcing, measurements of oscillating systems may be degraded by background noise that arises from other physical processes. A number of solar processes contribute to the noise polluting measurements of solar p-modes. The properties of these noise sources have implications for the design of instruments and of analysis methods.

Solar noise in p-mode measurements comes mostly from magnetic activity (at long timescales) and from convective motions (at shorter timescales). In the velocity signal, models of the noise roughly follow a power law, with the power spectral density P (in m^2/s^2 per Hz) estimated by Jiménez *et al.* (1988) given

approximately by

$$P(\nu) = 10^7 \nu^{-1.6} . \quad (11)$$

(Note that a typical p-mode with amplitude 10 cm/s and linewidth 1 μHz has a power spectral density of about 10^4 m^2/s^2 per Hz. At p-mode frequencies, this is 2 or 3 orders of magnitude larger than the solar background.) The contrast between p-modes and the solar background is smaller in intensity than in velocity measurements, for two reasons. First, the amplitudes of the p-modes are relatively small in intensity because the radiative cooling time in the photosphere is short compared to the p-mode periods. Thus, though the p-modes are almost adiabatic when considered globally, in the solar atmosphere (where they are observed), they are almost isothermal, with only small associated intensity changes. Second, the solar background at p-mode frequencies comes mostly from thermal convection, the essential feature of which is temperature variation. The p-modes in irradiance data therefore appear superposed on a larger solar background than in velocity data, but they are nevertheless clearly visible.

1.3 Limits to Observable Accuracy

The noise background and the statistical nature of the p-mode driving processes set definite limits on one's ability to measure the properties of the p-modes. The things one would like to measure regarding an individual p-mode are its amplitude A, its frequency ν_0, and its frequency HWHM Γ. Assuming the mode spectrum to be represented by a Lorentzian line profile with additive noise, the expectation value of the observed power is

$$\langle P(\nu) \rangle = \epsilon + \frac{A}{[1 + (\nu - \nu_0)^2/\Gamma^2)]} . \quad (12)$$

The most interesting of the parameters is perhaps the line center frequency ν_0. For observing times T that are short compared with the mode decay time τ, the frequency precision that may be attained is roughly the reciprocal of the length of the observing interval:

$$(\delta\nu_0)_{rms} = \frac{1.}{T} . \quad (13)$$

This just reflects the scaling property of Fourier transforms (eq. 5), a property that is sometimes dignified with names like "the uncertainty principle." In this case, and as long as the power in the line is large enough so that the resonance is clearly visible above the background noise, neither the noise level nor the statistical variations in power due to stochastic excitation make much difference to the expected precision.

A more common case when analyzing solar p-modes is that in which $T \geq \tau$. In this case the resonance line is fully resolved, and the stochastic nature of the